



DEVELOPMENT AND TESTING OF A HIGH-SPEED
REAL-TIME KINEMATIC PRECISE DGPS POSITIONING
SYSTEM BETWEEN TWO AIRCRAFT

THESIS

Christopher J. Spinelli, Major, USAF

AFIT/GCS/ENG/06-12

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

DEVELOPMENT AND TESTING OF A HIGH-SPEED
REAL-TIME KINEMATIC PRECISE DGPS POSITIONING
SYSTEM BETWEEN TWO AIRCRAFT

THESIS

Presented to the Faculty
Department of Electrical and Computer Engineering
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
In Partial Fulfillment of the Requirements for the
Degree of Master of Science in Computer Science

Christopher J. Spinelli, B.S.C.S.
Major, USAF


September 2006

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

DEVELOPMENT AND TESTING OF A HIGH-SPEED
REAL-TIME KINEMATIC PRECISE DGPS POSITIONING
SYSTEM BETWEEN TWO AIRCRAFT

Christopher J. Spinelli, B.S.C.S.
Major, USAF

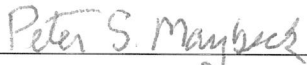
Approved:



Dr. John F. Raquet (Chairman)

30 AUG 06

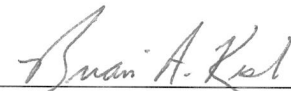
date



Dr. Peter S. Maybeck (Member)

30 August 2006

date



Major Brian A. Kish, Ph.D. (Member)

30 Aug 06

date

Abstract

This research involves the design, implementation, and testing of a high-speed, real-time kinematic, precise differential GPS positioning system for use in airborne applications such as automated aerial-refueling and close formation flying. Although many of the current ambiguity resolution techniques use the residuals from the least squares position estimation to determine the true ambiguity set, this thesis presents a novel approach to the ambiguity resolution problem, called the minimum indicator. Instead of assuming the ambiguity set with the lowest residuals is the true set, other special characteristics of the residuals are examined. This increases the confidence that the algorithm has selected the true ambiguity set. The end result was the first-ever successful in-flight demonstration of close formation flight, culminating in over 11 hours of close formation flying with a mean radial spherical error of 3.3 centimeters (0.108 feet). Other areas addressed include: the difference between “pre-fit” and “post-fit” residuals in the conditional probability calculation, the impact of a simplified dynamics model on system performance, the effect of widelane observables on the time-to-fix the correct double-difference ambiguities, and dynamically adjusting the time constant and standard deviation of relative acceleration states in the FOGMA model.

To my loving wife who is the real source of this thesis' production. Thanks for your patience, love, and wonderful ability to raise our sons.

Acknowledgements

First and foremost, I owe a large debt of gratitude to Dr. John Raquet. Without his help and guidance this thesis would not have been possible. His keen ability to determine the source of problems from limited information and how to correct them was invaluable. Also, his patience and approachable teaching style was greatly appreciated in light of my frequent visits to his office.

Two other individuals at AFIT also deserve recognition. John Schamus was the code work-horse for this thesis. He spent hours writing some of the more tedious code, and as much time helping me learn C++ and Linux. He was always there when I needed him, even as a “weekend warrior”. Don Smith, the hardware specialist, spent just as much time jumping through hoops to build two sets of “black boxes” and all the associated equipment. I guess we can forgive the Zener diode incident.

Next, I want to thank team “Lost Wingman” from TPS class 04B: Captain Benjamin “WIMS” George, Captain Adam “CHEWS” Faulkner, Captain York “BANG” Pasanen, Captain Scott “SULLY” Sullivan, and Mr. Bruce “SELF” Wilder. Your patience and hard work were instrumental in bringing the project together. “Could we get one more week for testing and coding here at AFIT?” Additionally, I want to thank team “No Gyro” from TPS class 05A: Major Steven “Burns” Ross, Major Aaron “Wheels” Mainstone, Captain Juanluis “Speedy” Velez, and Captain Elwood “Duck” Waddell Jr. The “No Gyro” team brought all of the pieces of the project together resulting in an awesome success.

The quality of this document would not have been possible without the thorough editing and great advice from my committee members, Dr. John Raquet, Major Brian Kish, Ph. D. and Dr. Peter Maybeck. Lastly, the numerous hours my lovely wife, also known as “The Comma Queen”, spent reviewing my work was greatly appreciated.

I would also like to thank my parents for their love and support while I was growing up and continue to receive. They instilled in me a strong work ethic and the value of education for as long as I can remember.

Last, but most definitely not least, I want to thank God for giving me the ability and drive to finish my AFIT experience. “And he said to me, My grace is sufficient for thee: for my strength is made perfect in weakness.” 2 Corinthians 12:9a

Christopher J. Spinelli

Table of Contents

	Page
Abstract	iv
Dedication	v
Acknowledgements	vi
List of Figures	xii
List of Tables	xxx
Notation	xxxii
List of Abbreviations	xxxiv
I. Introduction	1-1
1.1 Background	1-1
1.2 Problem Definition	1-5
1.3 Scope	1-5
1.4 Related Research	1-6
1.4.1 Kalman Filtering	1-6
1.4.2 Global Position System	1-6
1.5 Methodology	1-11
1.6 Thesis Overview	1-11
II. Theory	2-1
2.1 Overview	2-1
2.2 Global Position System	2-1
2.2.1 GPS Signal	2-1
2.2.2 GPS Measurements	2-3
2.2.3 Single Differencing	2-5
2.2.4 Double Differencing	2-7
2.2.5 Widelane Measurements	2-9
2.3 Carrier-Phase Ambiguity Resolution	2-10
2.3.1 LAMBDA	2-10
2.3.2 Ambiguity Set Determination	2-12
2.4 Kalman Filters	2-13
2.4.1 State Equations	2-13
2.4.2 State Transition Matrix	2-15
2.4.3 Measurement Model	2-15

	Page
2.4.4 State Estimates and Error Covariance Propagation	2-16
2.4.5 Measurement Incorporation	2-16
2.5 Multiple Model Adaptive Estimation	2-17
2.5.1 Conditional Probability Calculation	2-18
2.5.2 Overall MMAE State Estimate	2-20
2.6 Summary	2-21
III. Hardware Design and Algorithm Development	3-1
3.1 Overview	3-1
3.2 System Configuration	3-1
3.2.1 Data Collection	3-5
3.3 Overall Algorithm Structure	3-5
3.4 Floating Point DGPS Kalman Filter	3-5
3.4.1 Floating Point Filter Model Equations	3-9
3.4.2 Differential GPS Measurement Model	3-16
3.4.3 Discrete-Time Models	3-23
3.5 Candidate Ambiguity Set Generation	3-26
3.6 Multiple Model Adaptive Estimator Design	3-28
3.6.1 Elemental Filter Design	3-28
3.6.2 Conditional Probability Calculations	3-34
3.6.3 Elemental Filter Pruning	3-35
3.7 Minimum Indicator Design	3-37
3.7.1 Conditional Residuals Calculations	3-37
3.7.2 Sum-of-Squared Residual Description	3-41
3.7.3 Least Squares Fit Calculation	3-42
3.7.4 Minimum Indicator Calculation	3-44
3.8 Relative Position Update	3-46
3.9 Cycle Slips	3-47
3.9.1 Cycle Slip Detection	3-47
3.9.2 Cycle Slip Correction	3-51
3.9.3 MMAE and Minimum Indicator Cycle Slip Handling	3-51
3.9.4 Relative Position Update Cycle Slip Handling	3-53
3.10 Losing/Gaining Satellites	3-53
3.11 Changing Reference Satellite	3-54
3.12 Summary	3-57

	Page
IV. Results and Analysis	4-1
4.1 Overview	4-1
4.2 Test Background Information	4-1
4.2.1 Test Methodology	4-1
4.2.2 Data Sets	4-3
4.2.3 Data Reduction	4-4
4.3 Ground Tests	4-6
4.3.1 Case G1: Static Ground Test	4-7
4.3.2 Case G2: Dynamic Driving Test	4-15
4.3.3 Ground Test Summary	4-22
4.4 Flight Tests	4-22
4.4.1 Floating Filter Reset	4-25
4.4.2 True Set Not in Candidate Ambiguity Bank . .	4-33
4.4.3 Erroneous Ambiguity Set Local Minimum . . .	4-35
4.4.4 MMAE Outperforms Minimum Indicator	4-37
4.4.5 Twelve Visible Satellites	4-39
4.4.6 Multiple Satellite Dropouts	4-42
4.4.7 Real-time Timing	4-44
4.4.8 Flight Test Summary	4-45
4.5 Special Cases	4-46
4.5.1 “Pre-Fit” and “Post-Fit” Residual Comparison	4-46
4.5.2 Position-Only Filter Comparison	4-49
4.5.3 Widelane Observables Comparison	4-52
4.5.4 Dynamic T_a and σ_a^2 Comparison	4-57
4.5.5 Special Case Summary	4-60
4.6 Summary	4-60
V. Conclusions and Recommendations	5-1
5.1 Overview	5-1
5.2 Conclusions	5-1
5.3 Recommendations	5-5
Appendix A. Data Set Plots	A-1
A.1 Ground Tests	A-1
A.1.1 Case G1: Static Ground Test	A-1
A.1.2 Case G2: Dynamic Driving Test	A-16
A.1.3 Ground Summary	A-33
A.2 Flight Tests	A-33
A.2.1 Case F1.1: Flight Test 1, First Data Set	A-36
A.2.2 Case F1.2: Flight Test 1, Second Data Set . . .	A-51

	Page
A.2.3 Case F2.1: First Data Set Flight Test 2	A-69
A.2.4 Case F2.2: Flight Test 2, Second Data Set	A-81
A.2.5 Case F3.1: Flight Test 3, First Data Set	A-95
A.2.6 Case F3.2: Flight Test 3, Second Data Set	A-111
A.2.7 Case F4.1: Flight Test 4	A-127
A.2.8 Case F5.1: Flight Test 5, First Data Set	A-141
A.2.9 Case F5.2: Flight Test 5, Second Data Set	A-156
A.2.10 Case F6.1: First Flight of Flight Test 6, First Data Set	A-170
A.2.11 Case F6.2: First Flight of Flight Test 6, Second Data Set	A-185
A.2.12 Case F6.3: Second Flight of Flight Test 6, First Data Set	A-199
A.2.13 Case F6.4: Second Flight of Flight Test 6, Second Data Set	A-214
A.2.14 Case F7.1: Flight Test 7, First Data Set	A-229
A.2.15 Case F7.2: Flight Test 7, Second Data Set	A-242
A.2.16 Case F8.1: Flight Test 8, First Data Set	A-255
A.2.17 Case F8.2: Flight Test 8, Second Data Set	A-271
Bibliography	BIB-1
Vita	VITA-1

List of Figures

Figure		Page
1.1.	Typical GPS Accuracy	1-3
2.1.	Single Difference	2-6
2.2.	Double Difference	2-8
2.3.	Multiple Model Filter Algorithm	2-18
3.1.	Test Unit Installed on C-12	3-2
3.2.	Inside Back-plate of Test Unit	3-3
3.3.	Overall Message Flow	3-4
3.4.	Overall Algorithm Flow	3-6
3.5.	Algorithm Message Flow	3-8
3.6.	SOS Residuals Example	3-40
3.7.	SOS Residual Comparison	3-42
4.1.	Case G1: Floating Filter Position Errors	4-8
4.2.	Case G1: Satellite 26 Measurements	4-9
4.3.	Case G1: Visible Satellites	4-9
4.4.	Case G1: SOS Residuals	4-11
4.5.	Case G1: Conditional Probabilities for Selected MMAE Elemental Filters	4-11
4.6.	Case G1: Minimum Indicator and MMAE Position Difference	4-13
4.7.	Case G1: MMAE Position Error	4-14
4.8.	Case G1: Minimum Indicator Position Error	4-14
4.9.	Case G2: Floating Filter Position Errors	4-16
4.10.	Case G2: Satellite 6 Measurements	4-17
4.11.	Case G2: Visible Satellites	4-17
4.12.	Case G2: SOS Residuals	4-19

Figure		Page
4.13.	Case G2: Conditional Probabilities for Selected MMAE Elemental Filters	4-19
4.14.	Case G2: Conditional Probabilities for Selected MMAE Elemental Filters Expanded	4-20
4.15.	Case G2: MMAE Position Error	4-21
4.16.	Case G2: Minimum Indicator Position Error	4-21
4.17.	Every Epoch When the Aircraft Were Within 500 Feet	4-25
4.18.	Case F1.1: North, East, Down Relative Positions	4-26
4.19.	Case F1.1: Floating Filter Position Errors	4-27
4.20.	Case F1.1: Satellite 6 Measurements	4-28
4.21.	Case F1.1: Visible Satellites	4-29
4.22.	Case F1.1: SOS Residuals	4-30
4.23.	Case F1.1: Conditional Probabilities for Selected MMAE Elemental Filters	4-31
4.24.	Case F1.1: MMAE Position Error	4-32
4.25.	Case F1.1: Minimum Indicator Position Error	4-33
4.26.	Case F1.2: SOS Residuals	4-34
4.27.	Case F1.2: Conditional Probabilities for Selected MMAE Elemental Filters	4-35
4.28.	Case F2.2: SOS Residuals	4-36
4.29.	Case F2.2: Conditional Probabilities for Selected MMAE Elemental Filters	4-37
4.30.	Case F4.1: Minimum Indicator Position Error	4-38
4.31.	Case F4.1: MMAE Position Error	4-39
4.32.	Case F5.1: Visible Satellites	4-40
4.33.	Case F5.1: SOS Residuals	4-41
4.34.	Case F5.1: Conditional Probabilities for Selected MMAE Elemental Filters	4-41
4.35.	Case F8.2: Visible Satellites	4-42

Figure		Page
4.36.	Case F8.2: Minimum Indicator Position Error	4-43
4.37.	“Pre-fit” Residuals	4-46
4.38.	“Post-fit” Residuals	4-47
4.39.	“Pre-fit”/“Post-fit” Residual Case: Probability Comparison . .	4-48
4.40.	WL SOS Residuals	4-53
4.41.	L1 SOS Residuals	4-54
4.42.	WL SOS Residuals Expanded	4-55
4.43.	WL Conditional Probabilities for Selected MMAE Elemental Fil- ters	4-55
4.44.	L1 Conditional Probabilities for Selected MMAE Elemental Fil- ters	4-56
4.45.	Static T_a and σ_a^2 Floating Filter Velocity Errors	4-58
4.46.	Dynamic T_a and σ_a^2 Floating Filter Velocity Errors	4-58
4.47.	Static T_a and σ_a^2 Floating Filter Acceleration Errors	4-59
4.48.	Dynamic T_a and σ_a^2 Floating Filter Acceleration Errors	4-59
A.1.	Case G1: Floating Filter Position Errors	A-2
A.2.	Case G1: Floating Filter Velocity Errors	A-3
A.3.	Case G1: Floating Filter Acceleration Errors	A-3
A.4.	Case G1: Satellite 26 Measurements	A-4
A.5.	Case G1: Satellite 9 Measurements	A-5
A.6.	Case G1: Satellite 14 Measurements	A-5
A.7.	Case G1: Satellite 15 Measurements	A-6
A.8.	Case G1: Satellite 18 Measurements	A-6
A.9.	Case G1: Satellite 21 Measurements	A-7
A.10.	Case G1: Satellite 22 Measurements	A-7
A.11.	Case G1: Satellite 29 Measurements	A-8
A.12.	Case G1: Visible Satellites	A-8
A.13.	Case G1: SOS Residuals	A-10

Figure		Page
A.14.	Case G1: Conditional Probabilities for Selected MMAE Elemental Filters	A-11
A.15.	Case G1: EF[1] Probability Comparison	A-12
A.16.	Case G1: EF[2] Probability Comparison	A-13
A.17.	Case G1: EF[4] Probability Comparison	A-13
A.18.	Case G1: Minimum Indicator and MMAE Position Difference .	A-14
A.19.	Case G1: MMAE Position Error	A-15
A.20.	Case G1: Minimum Indicator Position Error	A-16
A.21.	Case G2: Floating Filter Position Errors	A-17
A.22.	Case G2: Floating Filter Velocity Errors	A-18
A.23.	Case G2: Floating Filter Acceleration Errors	A-18
A.24.	Case G2: Satellite 6 Measurements	A-19
A.25.	Case G2: Satellite 9 Measurements	A-20
A.26.	Case G2: Satellite 10 Measurements	A-20
A.27.	Case G2: Satellite 15 Measurements	A-21
A.28.	Case G2: Satellite 18 Measurements	A-21
A.29.	Case G2: Satellite 21 Measurements	A-22
A.30.	Case G2: Satellite 22 Measurements	A-22
A.31.	Case G2: Satellite 26 Measurements	A-23
A.32.	Case G2: Satellite 29 Measurements	A-23
A.33.	Case G2: Visible Satellites	A-24
A.34.	Case G2: SOS Residuals	A-26
A.35.	Case G2: Conditional Probabilities for Selected MMAE Elemental Filters	A-26
A.36.	Case G2: Conditional Probabilities for Selected MMAE Elemental Filters Expanded	A-27
A.37.	Case G2: EF[1] Probability Comparison	A-28
A.38.	Case G2: EF[2] Probability Comparison	A-28
A.39.	Case G2: EF[23] Probability Comparison	A-29

Figure		Page
A.40.	Case G2: EF[34] Probability Comparison	A-29
A.41.	Case G2: EF[36] Probability Comparison	A-30
A.42.	Case G2: EF[37] Probability Comparison	A-30
A.43.	Case G2: MMAE Position Error	A-31
A.44.	Case G2: Minimum Indicator Position Error	A-32
A.45.	Every Epoch When the Aircraft Were Within 500 Feet	A-36
A.46.	Case F1.1: North, East, Down Relative Positions	A-37
A.47.	Case F1.1: Floating Filter Position Errors	A-38
A.48.	Case F1.1: Floating Filter Velocity Errors	A-39
A.49.	Case F1.1: Floating Filter Acceleration Errors	A-39
A.50.	Case F1.1: Satellite 6 Measurements	A-40
A.51.	Case F1.1: Satellite 2 Measurements	A-41
A.52.	Case G2: Satellite 5 Measurements	A-41
A.53.	Case F1.1: Satellite 10 Measurements	A-42
A.54.	Case F1.1: Satellite 18 Measurements	A-42
A.55.	Case F1.1: Satellite 21 Measurements	A-43
A.56.	Case F1.1: Satellite 29 Measurements	A-43
A.57.	Case F1.1: Satellite 30 Measurements	A-44
A.58.	Case F1.1: Visible Satellites	A-45
A.59.	Case F1.1: SOS Residuals	A-46
A.60.	Case F1.1: Conditional Probabilities for Selected MMAE Elemental Filters	A-47
A.61.	Case F1.1: EF[1] Probability Comparison	A-48
A.62.	Case F1.1: EF[2] Probability Comparison	A-48
A.63.	Case F1.1: EF[23] Probability Comparison	A-49
A.64.	Case F1.1: MMAE Position Error	A-50
A.65.	Case F1.1: Minimum Indicator Position Error	A-51
A.66.	Case F1.2: North, East, Down Relative Positions	A-52

Figure		Page
A.67.	Case F1.2: Floating Filter Position Errors	A-54
A.68.	Case F1.2: Floating Filter Velocity Errors	A-54
A.69.	Case F1.2: Floating Filter Acceleration Errors	A-55
A.70.	Case F1.2: Satellite 15 Measurements	A-56
A.71.	Case F1.2: Satellite 2 Measurements	A-56
A.72.	Case G2: Satellite 5 Measurements	A-57
A.73.	Case F1.2: Satellite 10 Measurements	A-57
A.74.	Case F1.2: Satellite 16 Measurements	A-58
A.75.	Case F1.2: Satellite 18 Measurements	A-58
A.76.	Case F1.2: Satellite 21 Measurements	A-59
A.77.	Case F1.2: Satellite 22 Measurements	A-59
A.78.	Case F1.2: Satellite 26 Measurements	A-60
A.79.	Case F1.2: Satellite 29 Measurements	A-60
A.80.	Case F1.2: Satellite 30 Measurements	A-61
A.81.	Case F1.2: Visible Satellites	A-62
A.82.	Case F1.2: SOS Residuals	A-63
A.83.	Case F1.2: Conditional Probabilities for Selected MMAE Ele- mental Filters	A-63
A.84.	Case F1.2: EF[1] Probability Comparison	A-64
A.85.	Case F1.2: EF[21] Probability Comparison	A-65
A.86.	Case F1.2: EF[28] Probability Comparison	A-65
A.87.	Case F1.2: EF[37] Probability Comparison	A-66
A.88.	Case F1.2: EF[39] Probability Comparison	A-66
A.89.	Case F1.2: MMAE Position Error	A-67
A.90.	Case F1.2: Minimum Indicator Position Error	A-68
A.91.	Case F2.1: North, East, Down Relative Positions	A-69
A.92.	Case F2.1: Floating Filter Position Errors	A-71
A.93.	Case F2.1: Floating Filter Velocity Errors	A-71

Figure		Page
A.94.	Case F2.1: Floating Filter Acceleration Errors	A-72
A.95.	Case F2.1: Satellite 10 Measurements	A-73
A.96.	Case F2.1: Satellite 2 Measurements	A-73
A.97.	Case G2: Satellite 5 Measurements	A-74
A.98.	Case F2.1: Satellite 15 Measurements	A-74
A.99.	Case F2.1: Satellite 16 Measurements	A-75
A.100.	Case F2.1: Satellite 18 Measurements	A-75
A.101.	Case F2.1: Satellite 21 Measurements	A-76
A.102.	Case F2.1: Satellite 26 Measurements	A-76
A.103.	Case F2.1: Satellite 29 Measurements	A-77
A.104.	Case F2.1: Satellite 30 Measurements	A-77
A.105.	Case F2.1: Visible Satellites	A-78
A.106.	Case F2.1: MMAE Position Error	A-79
A.107.	Case F2.1: Minimum Indicator Position Error	A-81
A.108.	Case F2.2: North, East, Down Relative Positions	A-82
A.109.	Case F2.2: Floating Filter Position Errors	A-83
A.110.	Case F2.2: Floating Filter Velocity Errors	A-83
A.111.	Case F2.2: Floating Filter Acceleration Errors	A-84
A.112.	Case F2.2: Satellite 15 Measurements	A-85
A.113.	Case F2.2: Satellite 6 Measurements	A-85
A.114.	Case G2: Satellite 16 Measurements	A-86
A.115.	Case F2.2: Satellite 18 Measurements	A-86
A.116.	Case F2.2: Satellite 21 Measurements	A-87
A.117.	Case F2.2: Satellite 22 Measurements	A-87
A.118.	Case F2.2: Satellite 26 Measurements	A-88
A.119.	Case F2.2: Satellite 29 Measurements	A-88
A.120.	Case F2.2: Satellite 30 Measurements	A-89
A.121.	Case F2.2: Visible Satellites	A-90

Figure		Page
A.122.	Case F2.2: SOS Residuals	A-91
A.123.	Case F2.2: Conditional Probabilities for Selected MMAE Elemental Filters	A-91
A.124.	Case F2.2: EF[1] Probability Comparison	A-92
A.125.	Case F2.2: MMAE Position Error	A-93
A.126.	Case F2.2: Minimum Indicator Position Error	A-94
A.127.	Case F3.1: North, East, Down Relative Positions	A-96
A.128.	Case F3.1: Floating Filter Position Errors	A-97
A.129.	Case F3.1: Floating Filter Velocity Errors	A-97
A.130.	Case F3.1: Floating Filter Acceleration Errors	A-98
A.131.	Case F3.1: Satellite 16 Measurements	A-99
A.132.	Case F3.1: Satellite 1 Measurements	A-99
A.133.	Case G2: Satellite 3 Measurements	A-100
A.134.	Case F3.1: Satellite 4 Measurements	A-100
A.135.	Case F3.1: Satellite 13 Measurements	A-101
A.136.	Case F3.1: Satellite 20 Measurements	A-101
A.137.	Case F3.1: Satellite 23 Measurements	A-102
A.138.	Case F3.1: Satellite 24 Measurements	A-102
A.139.	Case F3.1: Satellite 25 Measurements	A-103
A.140.	Case F3.1: Satellite 27 Measurements	A-103
A.141.	Case F3.1: Visible Satellites	A-104
A.142.	Case F3.1: SOS Residuals	A-105
A.143.	Case F3.1: Conditional Probabilities for Selected MMAE Elemental Filters	A-106
A.144.	Case F3.1: EF[2] Probability Comparison	A-107
A.145.	Case F3.1: EF[4] Probability Comparison	A-107
A.146.	Case F3.1: EF[5] Probability Comparison	A-108
A.147.	Case F3.1: EF[6] Probability Comparison	A-108

Figure		Page
A.148.	Case F3.1: MMAE Position Error	A-109
A.149.	Case F3.1: Minimum Indicator Position Error	A-110
A.150.	Case F3.2: North, East, Down Relative Positions	A-111
A.151.	Case F3.2: Floating Filter Position Errors	A-113
A.152.	Case F3.2: Floating Filter Velocity Errors	A-113
A.153.	Case F3.2: Floating Filter Acceleration Errors	A-114
A.154.	Case F3.2: Satellite 16 Measurements	A-115
A.155.	Case F3.2: Satellite 3 Measurements	A-115
A.156.	Case G2: Satellite 8 Measurements	A-116
A.157.	Case F3.2: Satellite 13 Measurements	A-116
A.158.	Case F3.2: Satellite 19 Measurements	A-117
A.159.	Case F3.2: Satellite 20 Measurements	A-117
A.160.	Case F3.2: Satellite 23 Measurements	A-118
A.161.	Case F3.2: Satellite 27 Measurements	A-118
A.162.	Case F3.2: Visible Satellites	A-119
A.163.	Case F3.2: SOS Residuals	A-120
A.164.	Case F3.2: Conditional Probabilities for Selected MMAE Ele- mental Filters	A-120
A.165.	Case F3.2: EF[1] Probability Comparison	A-121
A.166.	Case F3.2: EF[1A] Probability Comparison	A-122
A.167.	Case F3.2: EF[1B] Probability Comparison	A-123
A.168.	Case F3.2: EF[2] Probability Comparison	A-123
A.169.	Case F3.2: EF[2A] Probability Comparison	A-124
A.170.	Case F3.2: EF[2B] Probability Comparison	A-124
A.171.	Case F3.2: EF[5] Probability Comparison	A-125
A.172.	Case F3.2: MMAE Position Error	A-125
A.173.	Case F3.2: Minimum Indicator Position Error	A-126
A.174.	Case F4.1: North, East, Down Relative Positions	A-127

Figure		Page
A.175.	Case F4.1: Floating Filter Position Errors	A-128
A.176.	Case F4.1: Floating Filter Velocity Errors	A-129
A.177.	Case F4.1: Floating Filter Acceleration Errors	A-129
A.178.	Case F4.1: Satellite 16 Measurements	A-130
A.179.	Case F4.1: Satellite 3 Measurements	A-131
A.180.	Case G2: Satellite 8 Measurements	A-131
A.181.	Case F4.1: Satellite 10 Measurements	A-132
A.182.	Case F4.1: Satellite 19 Measurements	A-132
A.183.	Case F4.1: Satellite 20 Measurements	A-133
A.184.	Case F4.1: Satellite 23 Measurements	A-133
A.185.	Case F4.1: Satellite 27 Measurements	A-134
A.186.	Case F4.1: Satellite 28 Measurements	A-134
A.187.	Case F4.1: Visible Satellites	A-135
A.188.	Case F4.1: SOS Residuals	A-136
A.189.	Case F4.1: Conditional Probabilities for Selected MMAE Elemental Filters	A-137
A.190.	Case F4.1: EF[1] Probability Comparison	A-138
A.191.	Case F4.1: EF[2] Probability Comparison	A-138
A.192.	Case F4.1: EF[3] Probability Comparison	A-139
A.193.	Case F4.1: MMAE Position Error	A-140
A.194.	Case F4.1: Minimum Indicator Position Error	A-141
A.195.	Case F5.1: North, East, Down Relative Positions	A-142
A.196.	Case F5.1: Floating Filter Position Errors	A-143
A.197.	Case F5.1: Floating Filter Velocity Errors	A-143
A.198.	Case F5.1: Floating Filter Acceleration Errors	A-144
A.199.	Case F5.1: Satellite 17 Measurements	A-145
A.200.	Case F5.1: Satellite 4 Measurements	A-145
A.201.	Case G2: Satellite 7 Measurements	A-146

Figure		Page
A.202.	Case G2: Satellite 8 Measurements	A-146
A.203.	Case F5.1: Satellite 9 Measurements	A-147
A.204.	Case F5.1: Satellite 11 Measurements	A-147
A.205.	Case F5.1: Satellite 24 Measurements	A-148
A.206.	Case F5.1: Satellite 26 Measurements	A-148
A.207.	Case F5.1: Satellite 27 Measurements	A-149
A.208.	Case F5.1: Satellite 28 Measurements	A-149
A.209.	Case F5.1: Satellite 29 Measurements	A-150
A.210.	Case F5.1: Satellite 31 Measurements	A-150
A.211.	Case F5.1: Visible Satellites	A-151
A.212.	Case F5.1: SOS Residuals	A-152
A.213.	Case F5.1: Conditional Probabilities for Selected MMAE Elemental Filters	A-153
A.214.	Case F5.1: EF[1] Probability Comparison	A-154
A.215.	Case F5.1: MMAE Position Error	A-155
A.216.	Case F5.1: Minimum Indicator Position Error	A-156
A.217.	Case F5.2: North, East, Down Relative Positions	A-157
A.218.	Case F5.2: Floating Filter Position Errors	A-158
A.219.	Case F5.2: Floating Filter Velocity Errors	A-158
A.220.	Case F5.2: Floating Filter Acceleration Errors	A-159
A.221.	Case F5.2: Satellite 8 Measurements	A-160
A.222.	Case F5.2: Satellite 2 Measurements	A-160
A.223.	Case F5.2: Satellite 4 Measurements	A-161
A.224.	Case G2: Satellite 7 Measurements	A-161
A.225.	Case F5.2: Satellite 9 Measurements	A-162
A.226.	Case F5.2: Satellite 24 Measurements	A-162
A.227.	Case F5.2: Satellite 26 Measurements	A-163
A.228.	Case F5.2: Satellite 28 Measurements	A-163

Figure		Page
A.229.	Case F5.2: Satellite 29 Measurements	A-164
A.230.	Case F5.2: Satellite 31 Measurements	A-164
A.231.	Case F5.2: Visible Satellites	A-165
A.232.	Case F5.2: SOS Residuals	A-166
A.233.	Case F5.2: Conditional Probabilities for Selected MMAE Elemental Filters	A-167
A.234.	Case F5.2: EF[1] Probability Comparison	A-168
A.235.	Case F5.2: MMAE Position Error	A-169
A.236.	Case F5.2: Minimum Indicator Position Error	A-170
A.237.	Case F6.1: North, East, Down Relative Positions	A-171
A.238.	Case F6.1: Floating Filter Position Errors	A-172
A.239.	Case F6.1: Floating Filter Velocity Errors	A-172
A.240.	Case F6.1: Floating Filter Acceleration Errors	A-173
A.241.	Case F6.1: Satellite 16 Measurements	A-174
A.242.	Case F6.1: Satellite 3 Measurements	A-174
A.243.	Case F6.1: Satellite 4 Measurements	A-175
A.244.	Case G2: Satellite 8 Measurements	A-175
A.245.	Case F6.1: Satellite 13 Measurements	A-176
A.246.	Case F6.1: Satellite 19 Measurements	A-176
A.247.	Case F6.1: Satellite 20 Measurements	A-177
A.248.	Case F6.1: Satellite 24 Measurements	A-177
A.249.	Case F6.1: Satellite 27 Measurements	A-178
A.250.	Case F6.1: Visible Satellites	A-179
A.251.	Case F6.1: SOS Residuals	A-180
A.252.	Case F6.1: Conditional Probabilities for Selected MMAE Elemental Filters	A-180
A.253.	Case F6.1: EF[1] Probability Comparison	A-181
A.254.	Case F6.1: EF[2] Probability Comparison	A-182

Figure		Page
A.255.	Case F6.1: EF[18] Probability Comparison	A-182
A.256.	Case F6.1: MMAE Position Error	A-183
A.257.	Case F6.1: Minimum Indicator Position Error	A-184
A.258.	Case F6.2: North, East, Down Relative Positions	A-185
A.259.	Case F6.2: Floating Filter Position Errors	A-186
A.260.	Case F6.2: Floating Filter Velocity Errors	A-187
A.261.	Case F6.2: Floating Filter Acceleration Errors	A-187
A.262.	Case F6.2: Satellite 16 Measurements	A-188
A.263.	Case F6.2: Satellite 3 Measurements	A-189
A.264.	Case G2: Satellite 8 Measurements	A-189
A.265.	Case F6.2: Satellite 13 Measurements	A-190
A.266.	Case F6.2: Satellite 19 Measurements	A-190
A.267.	Case F6.2: Satellite 20 Measurements	A-191
A.268.	Case F6.2: Satellite 23 Measurements	A-191
A.269.	Case F6.2: Satellite 27 Measurements	A-192
A.270.	Case F6.2: Satellite 28 Measurements	A-192
A.271.	Case F6.2: Visible Satellites	A-193
A.272.	Case F6.2: SOS Residuals	A-194
A.273.	Case F6.2: Conditional Probabilities for Selected MMAE Ele- mental Filters	A-195
A.274.	Case F6.2: EF[1] Probability Comparison	A-196
A.275.	Case F6.2: EF[2] Probability Comparison	A-196
A.276.	Case F6.2: MMAE Position Error	A-197
A.277.	Case F6.2: Minimum Indicator Position Error	A-198
A.278.	Case F6.3: North, East, Down Relative Positions	A-199
A.279.	Case F6.3: Floating Filter Position Errors	A-200
A.280.	Case F6.3: Floating Filter Velocity Errors	A-201
A.281.	Case F6.3: Floating Filter Acceleration Errors	A-201

Figure		Page
A.282.	Case F6.3: Satellite 17 Measurements	A-202
A.283.	Case F6.3: Satellite 4 Measurements	A-203
A.284.	Case G2: Satellite 7 Measurements	A-203
A.285.	Case G2: Satellite 8 Measurements	A-204
A.286.	Case F6.3: Satellite 9 Measurements	A-204
A.287.	Case F6.3: Satellite 11 Measurements	A-205
A.288.	Case F6.3: Satellite 24 Measurements	A-205
A.289.	Case F6.3: Satellite 26 Measurements	A-206
A.290.	Case F6.3: Satellite 27 Measurements	A-206
A.291.	Case F6.3: Satellite 28 Measurements	A-207
A.292.	Case F6.3: Satellite 29 Measurements	A-207
A.293.	Case F6.3: Satellite 31 Measurements	A-208
A.294.	Case F6.3: Visible Satellites	A-209
A.295.	Case F6.3: SOS Residuals	A-210
A.296.	Case F6.3: Conditional Probabilities for Selected MMAE Elemental Filters	A-210
A.297.	Case F6.3: EF[1] Probability Comparison	A-211
A.298.	Case F6.3: MMAE Position Error	A-212
A.299.	Case F6.3: Minimum Indicator Position Error	A-213
A.300.	Case F6.4: North, East, Down Relative Positions	A-214
A.301.	Case F6.4: Floating Filter Position Errors	A-215
A.302.	Case F6.4: Floating Filter Velocity Errors	A-216
A.303.	Case F6.4: Floating Filter Acceleration Errors	A-216
A.304.	Case F6.4: Satellite 11 Measurements	A-217
A.305.	Case F6.4: Satellite 4 Measurements	A-218
A.306.	Case G2: Satellite 7 Measurements	A-218
A.307.	Case G2: Satellite 8 Measurements	A-219
A.308.	Case F6.4: Satellite 9 Measurements	A-219

Figure		Page
A.309.	Case F6.4: Satellite 17 Measurements	A-220
A.310.	Case F6.4: Satellite 24 Measurements	A-220
A.311.	Case F6.4: Satellite 26 Measurements	A-221
A.312.	Case F6.4: Satellite 27 Measurements	A-221
A.313.	Case F6.4: Satellite 28 Measurements	A-222
A.314.	Case F6.4: Satellite 29 Measurements	A-222
A.315.	Case F6.4: Satellite 31 Measurements	A-223
A.316.	Case F6.4: Visible Satellites	A-224
A.317.	Case F6.4: SOS Residuals	A-225
A.318.	Case F6.4: Conditional Probabilities for Selected MMAE Elemental Filters	A-225
A.319.	Case F6.4: EF[1] Probability Comparison	A-226
A.320.	Case F6.4: MMAE Position Error	A-227
A.321.	Case F6.4: Minimum Indicator Position Error	A-228
A.322.	Case F7.1: North, East, Down Relative Positions	A-229
A.323.	Case F7.1: Floating Filter Position Errors	A-230
A.324.	Case F7.1: Floating Filter Velocity Errors	A-231
A.325.	Case F7.1: Floating Filter Acceleration Errors	A-231
A.326.	Case F7.1: Satellite 16 Measurements	A-232
A.327.	Case F7.1: Satellite 3 Measurements	A-233
A.328.	Case G2: Satellite 8 Measurements	A-233
A.329.	Case F7.1: Satellite 13 Measurements	A-234
A.330.	Case F7.1: Satellite 19 Measurements	A-234
A.331.	Case F7.1: Satellite 20 Measurements	A-235
A.332.	Case F7.1: Satellite 23 Measurements	A-235
A.333.	Case F7.1: Satellite 27 Measurements	A-236
A.334.	Case F7.1: Satellite 28 Measurements	A-236
A.335.	Case F7.1: Visible Satellites	A-237

Figure		Page
A.336.	Case F7.1: SOS Residuals	A-238
A.337.	Case F7.1: Conditional Probabilities for Selected MMAE Elemental Filters	A-239
A.338.	Case F7.1: EF[1] Probability Comparison	A-240
A.339.	Case F7.1: MMAE Position Error	A-241
A.340.	Case F7.1: Minimum Indicator Position Error	A-242
A.341.	Case F7.2: North, East, Down Relative Positions	A-243
A.342.	Case F7.2: Floating Filter Position Errors	A-244
A.343.	Case F7.2: Floating Filter Velocity Errors	A-244
A.344.	Case F7.2: Floating Filter Acceleration Errors	A-245
A.345.	Case F7.2: Satellite 10 Measurements	A-246
A.346.	Case F7.2: Satellite 3 Measurements	A-246
A.347.	Case G2: Satellite 8 Measurements	A-247
A.348.	Case F7.2: Satellite 19 Measurements	A-247
A.349.	Case F7.2: Satellite 23 Measurements	A-248
A.350.	Case F7.2: Satellite 27 Measurements	A-248
A.351.	Case F7.2: Satellite 28 Measurements	A-249
A.352.	Case F7.2: Visible Satellites	A-250
A.353.	Case F7.2: SOS Residuals	A-251
A.354.	Case F7.2: Conditional Probabilities for Selected MMAE Elemental Filters	A-251
A.355.	Case F7.2: EF[1] Probability Comparison	A-252
A.356.	Case F7.2: EF[2] Probability Comparison	A-253
A.357.	Case F7.2: MMAE Position Error	A-254
A.358.	Case F7.2: Minimum Indicator Position Error	A-255
A.359.	Case F8.1: North, East, Down Relative Positions	A-256
A.360.	Case F8.1: Floating Filter Position Errors	A-257
A.361.	Case F8.1: Floating Filter Velocity Errors	A-257

Figure		Page
A.362.	Case F8.1: Floating Filter Acceleration Errors	A-258
A.363.	Case F8.1: Satellite 4 Measurements	A-259
A.364.	Case F8.1: Satellite 1 Measurements	A-259
A.365.	Case G2: Satellite 3 Measurements	A-260
A.366.	Case F8.1: Satellite 8 Measurements	A-260
A.367.	Case F8.1: Satellite 13 Measurements	A-261
A.368.	Case F8.1: Satellite 16 Measurements	A-261
A.369.	Case F8.1: Satellite 19 Measurements	A-262
A.370.	Case F8.1: Satellite 20 Measurements	A-262
A.371.	Case F8.1: Satellite 24 Measurements	A-263
A.372.	Case F8.1: Satellite 25 Measurements	A-263
A.373.	Case F8.1: Satellite 27 Measurements	A-264
A.374.	Case F8.1: Visible Satellites	A-265
A.375.	Case F8.1: SOS Residuals	A-266
A.376.	Case F8.1: Conditional Probabilities for Selected MMAE Elemental Filters	A-266
A.377.	Case F8.1: EF[1] Probability Comparison	A-267
A.378.	Case F8.1: EF[2] Probability Comparison	A-268
A.379.	Case F8.1: EF[9] Probability Comparison	A-268
A.380.	Case F8.1: MMAE Position Error	A-269
A.381.	Case F8.1: Minimum Indicator Position Error	A-270
A.382.	Case F8.2: North, East, Down Relative Positions	A-271
A.383.	Case F8.2: Floating Filter Position Errors	A-272
A.384.	Case F8.2: Floating Filter Velocity Errors	A-273
A.385.	Case F8.2: Floating Filter Acceleration Errors	A-273
A.386.	Case F8.2: Satellite 8 Measurements	A-274
A.387.	Case G2: Satellite 3 Measurements	A-275
A.388.	Case F8.2: Satellite 13 Measurements	A-275

Figure		Page
A.389.	Case F8.2: Satellite 16 Measurements	A-276
A.390.	Case F8.2: Satellite 19 Measurements	A-276
A.391.	Case F8.2: Satellite 20 Measurements	A-277
A.392.	Case F8.2: Satellite 23 Measurements	A-277
A.393.	Case F8.2: Satellite 27 Measurements	A-278
A.394.	Case F8.2: Visible Satellites	A-279
A.395.	Case F8.2: SOS Residuals	A-280
A.396.	Case F8.2: Conditional Probabilities for Selected MMAE Ele- mental Filters	A-280
A.397.	Case F8.2: EF[1] Probability Comparison	A-281
A.398.	Case F8.2: EF[2] Probability Comparison	A-282
A.399.	Case F8.2: MMAE Position Error	A-282
A.400.	Case F8.2: Minimum Indicator Position Error	A-283

List of Tables

Table	Page
1.1. Typical GPS Accuracy	1-2
2.1. GPS Signal Structure	2-2
3.1. Javad® GPS Receiver Configuration	3-3
3.2. Dynamic T_a and σ_a^2 Progression	3-11
3.3. Static Floating Point Noise Values	3-14
3.4. Initial Floating Point Filter Covariance Values	3-16
3.5. Floating Point Filter Covariance Error Values	3-24
3.6. Doppler Change Caused by Clock Shift	3-50
4.1. Ground Maneuver Set for Relative GPS Position Testing	4-2
4.2. C-12C Aircraft Maneuver Set for Relative GPS Position Testing	4-3
4.3. Data Set Summary	4-5
4.4. Case G1: Position Error Summary (m)	4-15
4.5. Case G2: Position Error Summary (m)	4-22
4.6. Satellite Summary for Flight Testing	4-23
4.7. Flight Test Position Error Summary (m)	4-24
4.8. Case F1.1: Position Error Summary (m)	4-32
4.9. Velocity/Acceleration vs. Position-only	4-51
4.10. Widelane and L1 Observables Summary	4-56
A.1. Case G1: Position Error Summary (m)	A-15
A.2. Case G2: Position Error Summary (m)	A-32
A.3. Satellite Summary for Flight Testing	A-34
A.4. Flight Test Position Error Summary (m)	A-34
A.5. Case F1.1: Position Error Summary (m)	A-51

Table		Page
A.6.	Case F1.2: Position Error Summary (m)	A-69
A.7.	Case F2.1: Position Error Summary (m)	A-80
A.8.	Case F2.2: Position Error Summary (m)	A-94
A.9.	Case F3.1: Position Error Summary (m)	A-110
A.10.	Case F3.2: Position Error Summary (m)	A-126
A.11.	Case F4.1: Position Error Summary (m)	A-140
A.12.	Case F5.1: Position Error Summary (m)	A-155
A.13.	Case F5.2: Position Error Summary (m)	A-169
A.14.	Case F6.1: Position Error Summary (m)	A-184
A.15.	Case F6.2: Position Error Summary (m)	A-198
A.16.	Case F6.3: Position Error Summary (m)	A-213
A.17.	Case F6.4: Position Error Summary (m)	A-228
A.18.	Case F7.1: Position Error Summary (m)	A-241
A.19.	Case F7.2: Position Error Summary (m)	A-254
A.20.	Case F8.1: Position Error Summary (m)	A-270
A.21.	Case F8.2: Position Error Summary (m)	A-283

Notation

Scalars, Vectors, Matrices

Scalars - Denoted by upper or lower case letters in italic type, as in the scalar x .

Vectors - Denoted by lower case letters in boldface type, as in the vector \mathbf{x} made up of components x_i (always expressed as column vectors).

Matrices - Denoted by upper case letters in boldface type, as in the matrix \mathbf{A} made up of elements $A_{i,j}$ (i^{th} row, j^{th} column).

Prefixes

$E\{\cdot\}$ - Expected Value of $\{\cdot\}$

$\Delta(\cdot)$ - Incremental change in (\cdot) over a period of time, or
Single difference value of (\cdot) in GPS applications

$\delta(\cdot)$ - Error in the estimated, computed or measured (\cdot) w.r.t the true (\cdot)

$\Delta\nabla(\cdot)$ - Double difference value of (\cdot)

Overscores

$\hat{\cdot}$ - Estimated, Computed or Measured Value (not true value)

$\dot{\cdot}$ - Time derivative

$\bar{\cdot}$ - Mean value

Right Subscripts

- $_0$ - Initial, as in initial conditions
- $_{1,2}$ - Counting indices (usually used as sub-subscripts)
- $_{amb}$ - Ambiguity as in carrier-phase ambiguity
- $_d$ - Discrete-Time
- $_f$ - Filter Model
- $_{i,j}$ - Sample Time Index
- $_k$ - Elemental Filter Index
- $_{L1}$ - GPS L1 Frequency
- $_{L2}$ - GPS L2 Frequency
- $_{WL}$ - Widelane Observable
- $_n$ - Nominal (as in trajectory), or GPS Satellite number n
- $_{sv}$ - Satellite Vehicle
- $_u$ - User GPS Receiver
- $_{x,y}$ - Generic GPS Receiver index (Receiver x , Receiver y)

Right Superscripts

- $^{-1}$ - Matrix Inverse
- a,b - Generic GPS Satellite index (Satellite a , Satellite b)
- T - Vector or Matrix Transpose

List of Abbreviations

Abbreviation		Page
UAV	Unmanned Aerial Vehicle	1-1
TACAN	Tactical Air Navigation	1-2
DGPS	Differential GPS	1-2
WAAS	Wide Area Augmentation System	1-4
LAMBDA	Least-squares AMbiguity Decorrelation Adjustment	1-5
SOS	Sum-of-Squared	1-5
AFB	Air Force Base	1-6
SPS	Standard Positioning Service	1-7
PPS	Precise Positioning Service	1-7
ILS	Instrument Landing System	1-7
AFM	Ambiguity Function Method	1-9
LSAST	Least Squares Ambiguity Search Technique	1-9
FASF	Fast Ambiguity Search Filter	1-9
NASA	National Aeronautics and Space Administration	1-10
NM	Nautical Mile	1-10
EKF	Extended Kalman Filter	1-11
MMAE	Multiple Model Adaptive Estimation	1-11
PRN	Pseudo-Random Noise	2-1
SV	Satellite Vehicles	2-1
C/A	Course/Acquisition	2-1
P	Precision	2-1
BPSK	Binary Phase Shift Key	2-2
WL	Widelane	2-9
EF	Elemental Filters	2-17
RTK	Real-Time Kinematic	3-1

Abbreviation		Page
EAFB	Edwards Air Force Base	3-1
INU	Inertial Navigation Unit	3-2
GAINR	GPS Aided Inertial Navigation Reference	3-4
ECEF	Earth Centered Earth Fixed	3-9
FOGMA	First Order Gauss Markov Acceleration	3-10
TA	Total Acceleration	3-11
AFIT	Air Force Institute of Technology	3-41
PLL	Phase-Lock Loop	3-47
DLL	Delay-Lock Loop	3-47
WPAFB	Wright-Patterson Air Force Base	4-1
TPS	Test Pilot School	4-2
TMP	Test Management Project	4-2
MEMS	Microelectromechanical Systems	4-2
IMU	Inertial Measurement Unit	4-2
SLUF	Straight and Level Unaccelerated Flight	4-2
SHSS	Steady Heading Side Slips	4-2
KIAS	Knots Indicated Airspeed	4-3
PA	Pressure Altitude	4-3
DD	Double-Difference	4-6
DRMS	Distance Root Mean Square	4-7
MRSE	Mean Radial Spherical Error	4-7
RTB	Return To Base	4-38
CRC	Cyclic Redundancy Check	4-45

DEVELOPMENT AND TESTING OF A HIGH-SPEED REAL-TIME KINEMATIC PRECISE DGPS POSITIONING SYSTEM BETWEEN TWO AIRCRAFT

I. Introduction

1.1 Background

As the 21st century begins, computers have become an instrumental part of the human experience. From tiny microchips inside car keys, to huge super computers like IBM's Blue Gene/L, the uses for computer processing grows daily. In the world of aviation, microprocessors have revolutionized the industry both in capability and safety. In the same manner, GPS has ushered in a new era of navigation technologies. Complex avionic systems are becoming more common, even on light civilian aircraft. Technology has reached the point that even the pilot is being removed from the cockpit. Unmanned aerial vehicles (UAVs) represent the future for many civil and military airborne applications. UAVs such as the Predator and Global Hawk are already actively used in military applications, including wartime operations.

The removal of the pilot and other support personnel from an aircraft leaves room for items such as new hardware or more internal fuel. Additionally, human biological constraints are removed allowing for extended airborne operations. Sortie time constraints are now based on consumables such as engine lubrication and fuel. Currently, UAVs are not capable of aerial refueling. In addition, they normally operate single ship. Aerial refueling capability for UAVs would greatly increase their effectiveness and operational range. Also, increased firepower and capability are possible with formations of UAVs.

Both close formation flying and aerial refueling require precise flying. At the time of this thesis, only limited testing had been done on UAV formation flying, and almost no testing of autonomous aerial refueling. Most techniques for automated

precise formation flying involve determining the relative position between the aircraft and then building a flight controller around the position estimation. However, other techniques such as video recognition are currently under development [5].

To fly in a stable relative position, a flight controller needs precise, real-time position information at a high data rate. There are numerous ways the relative position can be determined. Some possible techniques include: laser ranging, video imaging, air-to-air tactical air navigation (TACAN), and differential GPS (DGPS) systems. This thesis involves a DGPS approach to the precise relative position problem.

DGPS is not concerned with absolute positioning, but the relative position between two or more receivers. DGPS is possible since many of the errors on the GPS measurements are spatially correlated. In other words, signals from GPS satellites experience similar atmospheric effects, especially when the receivers are in close proximity (within approximately 5 kilometers). DGPS offers significantly better performance than stand-alone GPS solutions. The typical accuracy of various GPS applications are shown in Table 1.1 and Figure 1.1 [41].

Table 1.1: Typical GPS Accuracy

	Mode	Approximate Horizontal Accuracy (m)
Autonomous	Civilian Receiver	6-10
	Military Receiver	3-6
Differential	Code DGPS	1-2
	Floating Point Carrier DGPS	0.2-0.5
	Fixed Integer Carrier DGPS	0.01-0.02

DGPS systems provide two types of solutions: code-based and carrier-based solutions. Code-based DGPS solutions rely on the pseudorange measurement generated from the GPS receiver. The pseudorange measurement is derived by determining the time difference between when the GPS signal is transmitted from the satellite and when the user receives the signal. GPS receivers perform the time difference by generating replica codes for known visible satellites, and then attempt to match the

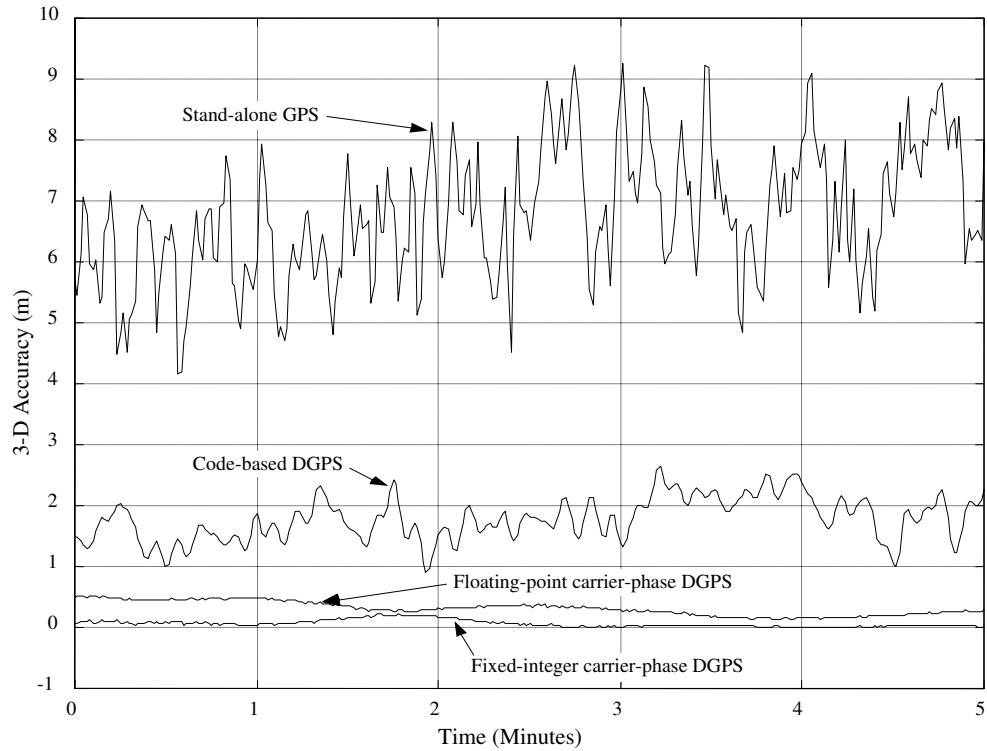


Figure 1.1: Typical GPS Accuracy [20]

incoming satellite codes with the receiver generated code using autocorrelation techniques. Once the codes are matched, the receiver knows the time the signal was sent and the time the signal was received. Since the GPS signal propagates through the atmosphere at the speed of light, a measurement of the distance between the receiver and satellite can be computed. The calculated range would be the actual distance between the receiver and satellite if:

- The receiver and satellite clocks were perfectly synchronized.
- There were no errors from the atmosphere.
- There were no errors from multipath.
- There was no noise from the internal components, wiring, and antennas of the satellite and receiver.

Unfortunately, all of those errors do exist, and the pseudorange is only an estimate of the true range. Code DGPS techniques attempt to remove/reduce the errors on

the pseudorange measurements. The Wide Area Augmentation System (WAAS) is a good example of a code-based DGPS system in widespread use today [41].

Code-based DGPS solutions offer an increased reduction in position errors, but they are still not accurate enough for applications such as precise formation flying between two aircraft. Safety requirements drive the need for centimeter-level or better position accuracy. Carrier-phase DGPS techniques can attain the needed position accuracy. Like code DGPS, carrier-phase DGPS also removes/reduces errors in GPS signal, but with greater effect.

There are two primary differences between code and carrier-phase DGPS. The first difference is that the carrier-phase measurement is much more precise than the code measurements, because it is based upon precisely tracking the phase of the GPS carrier signal. The second difference is the integer ambiguity bias on the carrier-phase measurements. Unlike the code measurements, which are an unambiguous range measurement, the carrier-phase measurements only indicate the current location within a cycle. There are an unknown number of integer cycles between the receiver and satellite, plus the partial cycle that the receiver is already measuring. To make the carrier-phase measurements unambiguous in range, the ambiguities must be estimated or resolved.

Estimates of the ambiguities are usually obtained through a Kalman filter [32] or simple code carrier differencing technique [47]. When the ambiguities are estimated in this manner, it is called a floating-point solution. Although the ambiguity estimates are not integers, they still allow for an unambiguous range measurement from the phase observables. Floating-point solutions offer a smaller variance and more accurate position solution than their code counter parts. However, the most accurate position solutions come when the ambiguities are fixed integers.

Normally, the ambiguity estimates from a Kalman filter, along with their covariance, are used to help resolve the integer ambiguities. Another technique for resolving the integer ambiguities is simply rounding the floating-point ambiguity estimates. The

rounding technique requires a wait time for the filter to converge properly. In either case, once the ambiguities are fixed, centimeter-level position accuracy is possible (assuming the correct ambiguity set is chosen).

There are generally two steps in the determination of the fixed-integer ambiguities. The first step is to construct a search space of candidate ambiguity sets. Next, the selection of the true, or most correct, ambiguity set is performed. There are many techniques for both steps of the integer ambiguity resolution process. This thesis used the Least-squares AMbiguity Decorrelation Adjustment (LAMBDA) method, developed by Teunissen, as an efficient process for performing the first step [53]. Determination of the true ambiguity set was the major thrust of this thesis. A new technique involving the special characteristics of the sum-of-squared (SOS) residuals from the least squares position estimation was developed.

1.2 Problem Definition

The primary goal of this thesis was the development and testing of a high-speed, real-time kinematic carrier-phase DGPS system for precise positioning between two aircraft, for use in applications such as aerial refueling and close formation flying. The system under development needed to be robust enough to deal with satellite masking, cycle slips, and datalink transmission errors. For safety, the system had to report the estimated accuracy to the pilot using an easy-to-understand scale. Next, the algorithm/datalink had to be fast enough to deliver the precise relative position at regular 20 Hertz intervals for use with a derivative-type flight controller. Finally, the system had to provide a centimeter-level or better relative position solution consistently during critical phases of flight, defined as when the aircraft were within 500 feet of each other.

1.3 Scope

The system was coded in C++ using the KDevelop environment in SUSE 9.2 [28, 42]. Data reduction on the output files was completed using MATLAB[®] [37]. All de-

velopment and testing used actual GPS measurements from Javad® JNS100 receivers (except the widelane example, which came from Ashtech GPS Z-Surveyors) [35]. Truth source data for all kinematic data sets, including the relative position, relative velocity, relative acceleration, and ambiguities, came from the GrafMov software suite developed by Waypoint Consulting, Inc. [36]. The ground tests were conducted at Wright-Patterson Air Force Base (AFB), OH from Jan 2005 to April 2005, and included both static tests and kinematic driving tests. Flight testing was done at Edwards AFB, CA in April and October 2005. April's flight tests were performed on two Air Force C-12s to verify the datalink functionality, and the algorithm's ability to provide the required centimeter-level relative position data. The April flights were a risk reduction step for the October flight tests, which included a flight controller wrapped around the precise relative position. During the October flight tests, a Calspan Learjet was used as the wing aircraft. It had a variable in-flight stability system, which allowed the input of a specialized set of flight control laws. The same Air Force C-12 from the April flight tests was used as the lead aircraft during the October flight tests. The results from selected ground and flight tests are presented in Chapter IV. Appendix A contains expanded results from all of the ground and flight testing.

1.4 Related Research

There has been extensive research in the areas of Kalman filtering and DGPS. Commercial industry has driven much of the research in both arenas, but academia laid the foundation for many of the techniques in use today.

1.4.1 Kalman Filtering. The Kalman filter was an important aspect of this thesis, but not the main thrust. An excellent background of Kalman filtering with related references can be found in [20].

1.4.2 Global Position System. The Global Position System consists of a baseline constellation of 24 satellites approximately 20,000 kilometers above the earth.

The system gives the user segment an estimation of the receiver's position and velocity. Additionally, precise timing is possible since each satellite has an atomic clock. Two kinds of service are currently offered by GPS: Standard Positioning Service (SPS) and Precise Positioning Service (PPS). SPS is generally reserved for civilian users. It offers an accuracy of approximately 6-10 meters, and was the service used during this thesis. PPS is only available to military receivers with special encryption keys, and has an accuracy of 3-6 meters. However, increased accuracy is possible using differential techniques.

1.4.2.1 Code DGPS. Code-based DGPS solutions are in widespread use today because of their relative ease. Code DGPS systems difference the pseudorange measurements from two or more receivers to produce a position accuracy of 1-2 meters. Since the code measurements are unambiguous in range, they can be used directly from the GPS receiver without any special processing. However, the code measurements are noisier than the carrier-phase measurements. Thus, the code measurements were not used to produce the final precise relative position output in this thesis.

Reliable and robust code-base DGPS systems first appeared during the early 1990s. A few of the initial code DGPS systems were used in farming applications with limited success [30, 46]. The tedious nature of driving huge farm machinery, precisely over long periods of time, during inclement weather or dusty conditions, drove the desire for automation. Most farm machinery travel at relatively slow speeds compared with highway traffic. However, car manufactures were investigating the use of DGPS for driving automation at higher speeds during the same time period [9].

Airports were also looking for inventive ways to update aging instrument landing systems (ILSs) with new low cost, low maintenance solutions. DGPS offered an elegant solution to this problem. Initial testing of code-base DGPS systems for precision approaches began as early as 1986 using a surveyed location on the airfield for the reference receiver [12]. As the systems matured, they proved reliable and accurate

enough to meet ILS Category 1 requirements both laterally and vertically [21]. Today GPS approaches can be found at airports around the country.

The United States Coast Guard was one of the major proponents for DGPS systems. Due to the Coast Guard's long history and expertise in the area of navigation, in 1989 it was tasked by the Department of Transportation to be the lead civil GPS interface. After taking over that role, the agency has help propel GPS navigation technologies. In 1999, a code-based differential system for the entire U.S. coastline, including the Great Lakes and major rivers was declared operational [52]. Properly equipped shipping vessels could navigate into port with an accuracy of 1-3 meters. These humble beginnings started a GPS expansion that would revolutionize navigation.

1.4.2.2 Carrier-phase DGPS. To unlock the true accuracy potential of DGPS positioning, the carrier-phase measurements must be utilized. Unfortunately, the carrier-phase measurements are ambiguous in range, and require substantial processing to resolve the integer ambiguities. Once the ambiguities are fixed though, relative position estimates of 0.01-0.02 meters are possible.

In the mid 1990s, as ambiguity resolution techniques were maturing, carrier-based DGPS solution became available. Initially, precise carrier-phase systems were very expensive (tens of thousands of dollars), and required very long initialization times (hours). As with the code-based DGPS system, farming was one of the largest user segments of the new carrier-phase differential systems [43]. However, instead of the DGPS system only being used to control the farm machinery, farmers also used the precise positioning to create topographic maps for the targeted placement of fertilizer [15]. The uses for carrier-phase DGPS has continued to grow, and the process of resolving the carrier-phase ambiguities has been a major source of research since GPS was first deployed.

1.4.2.3 Ambiguity Resolution. The early development of ambiguity resolution consisted of position-based approaches such as the Ambiguity Function Method (AFM) [8,31,49]. Position-based approaches were computationally expensive and time consuming. Generally, position-based approaches are no longer used.

In 1990, Hatch developed a new ambiguity resolution method called the Least Squares Ambiguity Search Technique (LSAST) [19]. After determining a bank of candidate ambiguity sets, a least squares position estimate of each ambiguity set was computed. The set with the lowest residuals from the least squares estimate was considered the true ambiguity set. While the LSAST technique was more efficient than the AFM, it was still computationally expensive, and is also not commonly used.

More advanced ambiguity-centered techniques, based on ambiguity estimates and their covariance, were developed in the mid 1990s. Two of these techniques were the Least-squares AMbiguity Decorrelation Adjustment (LAMBDA) [53] and Fast Ambiguity Search Filter (FASF) [7]. The LAMBDA method was used in this thesis for ambiguity set generation. More detailed information on the LAMBDA method is presented in Section 2.3.

Later in 2001, an ambiguity resolution technique devised by Henderson was the first to use a multiple filter approach [20]. He showed that multiple filters could successfully resolve the carrier-phase ambiguities in widelane measurements with a baseline distance over 32 kilometers. The drawback to Henderson’s approach was the Bayesian estimator used for the final state estimates. Incorrect ambiguity sets induced errors in the relative position estimate. This thesis expanded on Henderson’s work and reduced the error caused by the Bayesian estimator by eliminating low probability sets from the estimation.

Although many of the current ambiguity resolution techniques use the residuals from the least squares position estimation to determine the true ambiguity set, this thesis presents a novel approach to the ambiguity resolution problem. Instead of as-

suming the ambiguity set with the lowest residuals is the true set, other characteristics of the residuals are examined. This increases the confidence that the algorithm has selected the true ambiguity set.

1.4.2.4 Flight Oriented DGPS Systems. Until recently, there has not been much research in the area of aircraft-to-aircraft DGPS systems since such systems are mostly confined to military applications. In 2002, the National Aeronautics and Space Administration (NASA) demonstrated a semi-autonomous formation flight between two F-18 Hornets [18]. The flight was designed as concept demonstration of autonomous station keeping. The aircraft were no closer than 56 feet during testing for safety considerations. Eventually, NASA intends on flying the aircraft closer together, with the ultimate goal of using the lead aircraft's wingtip vortices for drag reduction.

Another research project which investigated the drag reduction effect of the wingtip vortices was conducted by Osteros [45]. Unfortunately, required parameters from the lead aircraft were not available in real-time, and the tests were eventually flown single-ship. However, during the single-ship operations, a datalink connection was still used to maintain formation position based on a "ghost" flight-lead. Osteros ran into more problems during testing due to intermittent datalink operation. His lessons learned paid big dividends and resulted in the overall success of the October flight tests [51].

The first documented report of actual UAVs flying in formation occurred in August of 2004 [10]. Two Boeing X-45A aircraft were flown with a surrogate lead. However, the formation flight was limited to only 1 nautical mile (NM) ranges. This was a big step for the UAV community, but the aircraft were still not flown in close formation positions. This thesis builds on previous research, but eventually leads to the first documented report of two aircraft flying in autonomous, 6 degree of freedom close formation and simulating aerial refueling [13].

1.5 Methodology

The research in this thesis began with an in-depth study of Kalman filtering and DGPS techniques. Additionally, significant time was spent learning how to code in C++ and use the KDevelop environment. Next, the system hardware was built in conjunction with code development. The system was first tested statically in the laboratory with a known relative position. This verified the algorithm could handle the simplest case and operate in real-time. Dynamic driving tests followed, which stressed the system under kinematic conditions, verified the datalink operation to useful ranges, and determined optimal tuning parameters for the Kalman filter. Finally, flight testing was conducted with an operational system. Results are presented and recommendations are made in Chapters IV and V.

1.6 Thesis Overview

Chapter II lays the ground work for the information presented in future chapters. Included is a discussion of the Extended Kalman Filter (EKF) and Multiple Model Adaptive Estimation (MMAE). Basic GPS theory and other more involved subjects, such as differential GPS (DGPS), are presented next. Finally, the nature of the phase residuals from the EKF is presented.

Chapter III describes the hardware configuration, overall algorithm, system models, and techniques used throughout this thesis. First, an overall system description is given, including the hardware, followed by a more detailed discussion of each sub-system. Next, the Kalman filter states, equations, and other new techniques developed are given. Finally, areas of additional functionality are explained.

Chapter IV presents the results and analysis from ground and flight testing of the algorithm developed in Chapter III. First, background information to include test methodology and data set descriptions are presented. Ground test results follow for a static test and a dynamic driving test. Next, flight test results from April and October 2005 are shown. Finally, special cases such as “pre-fit” and “post-fit”

residual comparison, the floating filter in position-only mode, widelane measurements, and dynamic T_a and σ_a^2 are presented.

Chapter V summarizes the research effort, and presents conclusions and recommendation for follow-on research.

II. Theory

2.1 Overview

This chapter lays the groundwork for the information presented in future chapters. Included is a discussion of basic GPS theory and differential GPS (DGPS). Extended Kalman Filtering (EKF) and Multiple Model Adaptive Estimation (MMAE) are presented next. Much of the content for this chapter was adapted from [20, 22].

2.2 Global Position System

GPS has been around since the early 1980s, but was not declared operational until 1995 [41]. Since that time it has become one of the most widespread navigation technologies in use throughout the world. GPS has three major segments: the Space Segment, the Control Segment, and the User Segment [41]. The Space Segment consists of the GPS satellites themselves. The Control Segment is concerned with satellite management and estimation of satellite navigation parameters. Finally, the User Segment consists of everyone that uses a receiver to acquire the GPS signals.

2.2.1 GPS Signal. Currently, each GPS satellite transmits continuously on two carrier frequencies in the L-band, designated L1 (1575.42 MHz) and L2 (1227.60 MHz). The carrier frequencies are modulated by pseudo-random noise (PRN) codes. The PRN codes are sequences of ones and zeros that uniquely identify each satellite vehicle (SV). Since there are at least 24 satellites transmitting on the same two frequencies, the PRN codes are carefully generated to prevent interference with each other.

Each satellite transmits two different codes: a coarse/acquisition (C/A) code, intended for civil use, and the precision (P) code, for military use. The L1 frequency is modulated with both codes, but the L2 frequency is only modulated with the P-code. The C/A-code is modulated at a chipping rate of 1023 chips per millisecond, and it repeats every millisecond. The term chip rather than data bit is used for the C/A-code and P-code since no data is transferred through the codes. Because of the

C/A-code's high repeat rate, it is much easier to acquire than the P-code. The P-code has a chipping rate ten times that of the C/A-code, and only repeats every 7 days. Since 1994, the P-code has been encrypted with the Y-code, and this signal is often referred to as the P(Y)-code. The encryption is intended to prevent an adversary from transmitting a "copy-cat" P-code signal.

A 50-bit-per-second navigation message is also included on each frequency. The navigation message consists of satellite health information, ephemeris (satellite position and velocity), clock bias parameters, and an almanac giving reduced-precision ephemeris data on all satellites in the constellation [41]. The code signals and navigation message are first combined using modulo-2 arithmetic. The result is then modulated on the carrier frequency using binary phase shift key (BPSK) modulation. If the combined code/navigation message result is a zero, the carrier signal is left unchanged. If the result is a one, the carrier is multiplied by -1, creating a 180 degree phase shift in the carrier signal [41]. Table 2.1 gives a summary of the GPS signal structure.

Table 2.1: GPS Signal Structure

Carrier	Frequency	Wavelength	Modulation	Chipping Rate	Chip Length
L1	1575.42 MHz	19.03 cm	C/A-code	1.023 Mchips/sec	293.0 m
			P-code	10.23 Mchips/sec	29.3 m
			Nav Message	50 bits/sec	
L2	1227.60 MHz	24.42 cm	P-code	10.23 Mchips/sec	29.3 m
			Nav Message	50 bits/sec	

Receivers are either dual-frequency or single-frequency capable. If the receiver is dual-frequency capable, it can track both the L1 and L2 frequencies. Even though the P(Y)-code is encrypted, civil receiver manufacturers designed proprietary techniques to pseudo-track the L2 frequency [55]. However, the measurements from the L2 frequency are not as robust as a military receiver with the encryption keys. Single-frequency receivers can only track the L1 frequency. A majority of civil receivers are single-frequency. Dual-frequency GPS receivers are generally better because they

have the ability to form multiple measurement observables in order to reduce errors from the atmosphere, or increase the ability of resolving the integer ambiguities on the carrier-phase measurements.

The GPS signal is only a means to an end. The real power in GPS is its ability to provide accurate position, velocity, and timing information to users around the world. The following sections describe how precise positioning is accomplished.

2.2.2 GPS Measurements. There are two primary observables generated by GPS receivers — a pseudorange based on the code signal and a phase measurement based on the carrier-frequency. To obtain the code measurements, a receiver generates a replica code and matches the transmit time from the appropriate SV by using an autocorrelation process. The code measurements are unambiguous (i.e., they give a direct reading of the range between the user and the satellite). The phase measurements are an integration of the number of cycles of Doppler shift that have passed since the receiver phase loop started tracking a particular signal. The carrier-phase measurement is ambiguous since the receiver only knows the fractional portion of the initial phase, but there are an unknown number of integer cycles between the receiver and the satellite. This unknown bias in the phase measurements is referred to as the integer ambiguity. Each measurement is discussed in detail next.

2.2.2.1 Code Measurements. Although a code measurement is unambiguous, it still contains significant errors. These errors are the reason the code measurement is called a pseudorange (i.e., it contains both the true range plus the errors). The code measurement can be described by:

$$\rho = r + c(\delta t_u - \delta t_{sv}) + T + I + m + v \quad (2.1)$$

where

$\rho =$	pseudorange (m)
$r =$	true range between the satellite and receiver (m)
$c =$	speed of light (299792458 m/s)
$\delta t_u =$	user (receiver) clock error (sec)
$\delta t_{sv} =$	satellite clock error (sec)
$T =$	measurement delay due to troposphere (m)
$I =$	measurement delay due to ionosphere (m)
$m =$	measurement delay due to multipath (m)
$v =$	measurement noise (m)

Without any augmentation, single-point positioning based on the code measurements is generally within 10 meters [41]. Considering GPS satellites are on average 20,000 km above earth, this is an amazing feat. To get more precise positioning, phase measurements can be used.

2.2.2.2 Phase Measurements. The phase measurements are much more accurate than their code counterparts. The increased accuracy comes from a couple of sources. First, the multipath and noise errors on the phase measurements are generally smaller than the code measurements [48]. Also, the length of one phase cycle is three orders of magnitude smaller than the width of one code chip (0.2 m vs. 293.0 m). However, to utilize that accuracy, significantly more processing is required.

The phase measurement equation is similar to the code and is given as:

$$\phi = \frac{1}{\lambda} [r + c(\delta t_u - \delta t_{sv}) + T - I + m + v] + N \quad (2.2)$$

where r , c , δt_u , δt_{sv} , T , I , m , and v are defined in Equation (2.1), and

$\phi =$	phase measurement (cycles)
$\lambda =$	carrier wavelength (m/cycle)
$N =$	integer carrier-phase cycle ambiguity (cycles)

The phase measurement errors are similar to those from the code measurements. In fact, the tropospheric error, T , and both clock errors, δt_u and δt_{sv} , are the same values as the code equation. The only difference in the ionospheric error is the sign. The ionosphere advances the phase, but delays the code [29]. The phase multipath and noise use the same symbols, but their values will differ from the code (generally smaller).

The last term in Equation (2.2) is the integer ambiguity. As mentioned at the beginning of this section, the phase measurement is ambiguous in range. A GPS receiver is continuously integrating the Doppler from the time the phase loop locks onto the carrier signal. It “knows” how many cycles have past since the start of the tracking loop, but there are an initial N cycles between the satellite and receiver when the tracking loop started. N remains a constant bias unless a cycle slip or loss-of-lock occurs. To use the phase measurements for positioning, (i.e., make the phase a true range measurement) the integer ambiguity must be estimate or resolved. There are many methods for resolving the carrier-phase ambiguities, and some of the more advanced techniques can be found in [6, 7, 54, 56].

2.2.3 Single Differencing. Differencing techniques can be used to help reduce or eliminate the nuisance parameters from Equations (2.1) and (2.2). Differential GPS got its namesake from the differencing techniques. Single differencing is the first step in reducing GPS errors. Figure 2.1 shows the concept of single differencing.

Superscripts on the measurements signify satellite(s), while subscripts signify the receiver(s). The actual single difference is formed by:

$$\Delta\rho_{xy}^a = \rho_x^a - \rho_y^a \quad (2.3)$$

where

ρ_x^a = pseudorange between satellite a and receiver x (m)

ρ_y^a = pseudorange between satellite a and receiver y (m)

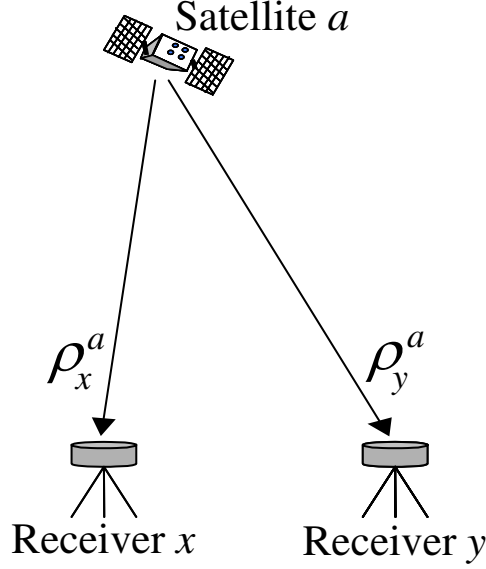


Figure 2.1: Single Difference [20]

Substituting the full code measurement specification from Equation (2.1) into Equation (2.3) gives:

$$\begin{aligned} \Delta\rho_{xy}^a &= [r_x^a + c(\delta t_{ux}^a - \delta t_{svx}^a) + T_x^a + I_x^a + m_x^a + v_x^a] \\ &\quad - [r_y^a + c(\delta t_{uy}^a - \delta t_{svy}^a) + T_y^a + I_y^a + m_y^a + v_y^a] \end{aligned} \quad (2.4)$$

Rearranging Equation (2.4) and collecting like terms gives:

$$\begin{aligned} \Delta\rho_{xy}^a &= (r_x^a - r_y^a) + [c(\delta t_{ux}^a - \delta t_{uy}^a) - c(\delta t_{svx}^a - \delta t_{svy}^a)] \\ &\quad + (T_x^a - T_y^a) + (I_x^a - I_y^a) + (m_x^a - m_y^a) + (v_x^a - v_y^a) \end{aligned} \quad (2.5)$$

The satellite clock error, δt_{sv}^a , completely cancels since the satellite transmitted to both receivers at the same time. Also, both the tropospheric and ionospheric errors are reduced. The closer the receivers are to each other, the more the atmospheric errors cancel, since the signals travel through the same piece of sky. There is a drawback to single differencing — the multipath and noise are amplified by a factor of $\sqrt{2}$ [47]. Now that the single difference equation is built, the final single difference

pseudorange equation is shown using the single difference notation from above:

$$\Delta\rho_{xy}^a = \Delta r_{xy}^a + \Delta c\delta t_{u_{xy}}^a + \Delta T_{xy}^a + \Delta I_{xy}^a + \Delta m_{xy}^a + \Delta v_{xy}^a \quad (2.6)$$

The same treatment is given to the phase measurements and results in the following equation:

$$\Delta\phi_{xy}^a = \frac{1}{\lambda} [\Delta r_{xy}^a + \Delta c\delta t_{u_{xy}}^a + \Delta T_{xy}^a - \Delta I_{xy}^a + \Delta m_{xy}^a + \Delta v_{xy}^a] + \Delta N_{xy}^a \quad (2.7)$$

The only differences between the single difference code and single difference phase equations are the additional ambiguity term in the phase equation, the units of each equation, the magnitude of the error terms, and the sign on the ionospheric error term. The new single difference ambiguity term, ΔN_{xy}^a , is still an integer, since subtraction was the only operation performed.

Single differencing reduces the error on the code and phase measurements, but an additional difference further eliminates nuisance terms without a large penalty. The additional difference is called double differencing.

2.2.4 Double Differencing. Double differencing is the difference between two single difference measurements. An example of how double differencing works is shown in Figure 2.2

The double difference observables are formed by:

$$\Delta\nabla\rho_{xy}^{ab} = \Delta\rho_{xy}^a - \Delta\rho_{xy}^b \quad (2.8)$$

where

$\Delta\rho_{xy}^a =$ single difference value between satellite a and receivers x and y (m)

$\Delta\rho_{xy}^b =$ single difference value between satellite b and receivers x and y (m)

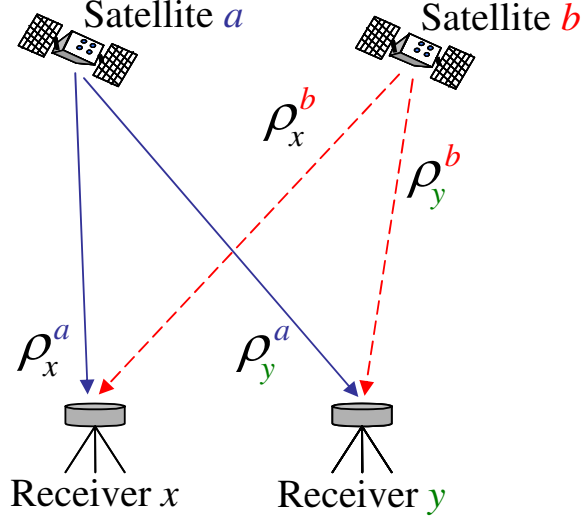


Figure 2.2: Double Difference [20]

Expanding Equation (2.8) and collecting like terms results in:

$$\begin{aligned} \Delta\nabla\rho_{xy}^{ab} &= \Delta r_{xy}^a + \Delta c\delta t_{u_{xy}}^a + \Delta T_{xy}^a + \Delta I_{xy}^a + \Delta m_{xy}^a + \Delta v_{xy}^a \\ &\quad - \Delta r_{xy}^b - \Delta c\delta t_{u_{xy}}^b - \Delta T_{xy}^b - \Delta I_{xy}^b - \Delta m_{xy}^b - \Delta v_{xy}^b \end{aligned} \quad (2.9)$$

The satellite clock errors were already removed during the single difference. The double difference will also remove the user clock errors, under the assumption that the measurements were received simultaneously at each receiver (i.e., $\delta t_{u_x}^a = \delta t_{u_x}^b$ and $\delta t_{u_y}^a = \delta t_{u_y}^b$). The tropospheric and ionospheric errors are still reduced. However, the elimination of the user clock error term was not without cost, because the error on the multipath and noise terms now increases by a factor of 2 [47]. Finally, the double difference pseudorange equation becomes:

$$\Delta\nabla\rho_{xy}^{ab} = \Delta\nabla r_{xy}^{ab} + \Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab} \quad (2.10)$$

The phase measurements are developed in the same manner and result in:

$$\Delta\nabla\phi_{xy}^{ab} = \frac{1}{\lambda} [\Delta\nabla r_{xy}^{ab} + \Delta\nabla T_{xy}^{ab} - \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab}] + \Delta\nabla N_{xy}^{ab} \quad (2.11)$$

Excellent discussions of the errors from Equations (2.10) and (2.11) can be found in [41, 48].

Double difference code and phase measurements from the L1 frequency were the primary observables used in this thesis. However, the benefits of using another measurement observable based on an L1-L2 combination, called widelane (WL), is shown in Chapter IV.

2.2.5 Widelane Measurements. The wavelength of the L1 and L2 signals are approximately 19 and 24 centimeters respectively. However, if a receiver is capable of tracking both the L1 and L2 signals, a new observable, called widelane (WL), is formed by the following equation:

$$\phi_{WL} = \phi_{L1} - \phi_{L2} \quad (2.12)$$

where

$$\begin{aligned} \phi_{WL} &= \text{Widelane phase observables} \\ \phi_{L1} &= \text{L1 phase observables} \\ \phi_{L2} &= \text{L2 phase observables} \end{aligned}$$

The widelane observable has the benefit of a much longer effective wavelength. The new WL wavelength is defined by:

$$\lambda_{WL} = \frac{\lambda_{L1}\lambda_{L2}}{\lambda_{L2} - \lambda_{L1}} \quad (2.13)$$

where

$$\begin{aligned} \lambda_{WL} &= \text{Widelane wavelength } (\approx 0.86192 \text{ m}) \\ \lambda_{L1} &= \text{L1 wavelength } (\approx 0.19029 \text{ m}) \\ \lambda_{L2} &= \text{L2 wavelength } (\approx 0.24421 \text{ m}) \end{aligned}$$

The longer wavelength makes determining the true integer ambiguity set more apparent than with just the pure carrier-phase alone. However, the penalty is in-

creased observation noise. The phase multipath, and noise on the position error, increase by a factor of approximately 6.4 [47]. A more detailed discussion of the costs and benefits of widelane measurements is in [48]. Although widelane measurements make the carrier-phase ambiguities more apparent, a method is still needed to construct the ambiguity search space and determine the true ambiguity set.

2.3 *Carrier-Phase Ambiguity Resolution*

To unlock the precise positioning with the carrier-phase measurements, the ambiguity term must be resolved. There are many approaches to ambiguity resolution. The early development of ambiguity resolution consisted of position-based approaches such as the Ambiguity Function Method (AFM) [8,31,49]. Position-based approaches were computationally expensive and time consuming. More advanced ambiguity-centered techniques, based on ambiguity estimates and their covariance, were developed in the mid 1990s. Two of these techniques were the Least-squares AMbiguity Decorrelation Adjustment (LAMBDA) [53] and Fast Ambiguity Search Filter (FASF) [7].

Generally, most carrier-phase resolution algorithms perform two steps. The first step is to determine the ambiguity search space (i.e., a set of candidate ambiguity sets that might be correct). The second step is to select the true ambiguity set (assuming it exists in the search space). The LAMBDA function developed by Teunissen [26,53] was used in this thesis to perform these functions.

2.3.1 LAMBDA. The algorithm used in this thesis came directly from [26]. A general overview of the LAMBDA method is provided as a reference for the reader. More in-depth discussions of the LAMBDA method are in [23–25, 27, 53, 54].

The ambiguity estimates from the floating filter contain a high degree of correlation which make ambiguity resolution in their native form difficult. To help reduce this correlation, Teunissen developed a Z-transform. After the ambiguities are trans-

formed, a much quicker and efficient search can be performed. To preserve the integer nature of the ambiguities, the transform must meet certain conditions:

1. Must be volume preserving (one-to-one relation).
2. Must reduce the product of ambiguity variances.
3. Must have integer elements.

The Z-transform is defined by:

$$\mathbf{z} = \mathbf{Z}^T \mathbf{x} \quad \hat{\mathbf{z}} = \mathbf{Z}^T \hat{\mathbf{x}} \quad \mathbf{Q}_{\hat{\mathbf{z}}} = \mathbf{Z}^T \mathbf{Q}_{\hat{\mathbf{x}}} \mathbf{Z} \quad (2.14)$$

where

$\mathbf{x}, \mathbf{z} =$ Untransformed and transformed ambiguities, respectively

$\mathbf{Z} =$ Transformation matrix

$\hat{\mathbf{x}}, \hat{\mathbf{z}} =$ Untransformed and transformed ambiguity estimates, respectively

$\mathbf{Q}_{\hat{\mathbf{x}}}, \mathbf{Q}_{\hat{\mathbf{z}}} =$ Untransformed and transformed covariance matrix, respectively

A simple example case is presented first. Assuming the ambiguity covariance matrix is given by:

$$\mathbf{Q}_{\hat{\mathbf{x}}} = \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{21} & \sigma_2^2 \end{bmatrix} \quad (2.15)$$

The Z-transform matrices would then be:

$$\mathbf{Z}_1^T = \begin{bmatrix} 1 & \text{int}(-\sigma_{12}\sigma_2^{-2}) \\ 0 & 1 \end{bmatrix} \quad (2.16)$$

or

$$\mathbf{Z}_2^T = \begin{bmatrix} 1 & 0 \\ \text{int}(-\sigma_{21}\sigma_1^{-2}) & 1 \end{bmatrix} \quad (2.17)$$

Either the upper (\mathbf{Z}_1^T) or the lower (\mathbf{Z}_2^T) triangular matrix can be used for the transformation. The “int(·)” in Equations (2.16) and (2.17) refers to rounding (·) to the nearest integer. This is done to preserve the integer nature of the ambiguities,

otherwise the transformation matrix would transform the ambiguities into non-integer quantities. Equations (2.16) and (2.17) worked for this simple case, but a more advanced method is needed for decorrelating more than two ambiguities.

The method used in the LAMBDA software is based on an **LDL** decomposition of $\mathbf{Q}_{\hat{\mathbf{x}}}$. A complete discussion of how the decomposition is performed can be found in [23]. After the decomposition is complete, the next step is to define the search space. The search space needs to be scaled so that it includes at least one point (usually two for validation purposes discussed later). The search space is described by:

$$(\hat{\mathbf{z}} - \mathbf{z})^T \mathbf{Q}_{\hat{\mathbf{z}}}^{-1} (\hat{\mathbf{z}} - \mathbf{z}) \leq \chi^2 \quad (2.18)$$

2.3.2 Ambiguity Set Determination. The search creates a collection of potential candidate ambiguity sets to consider. Now the problem becomes determining which candidate set is the true ambiguity set. To help determine the true ambiguity set, a weighted square norm of the position estimation from each candidate ambiguity set, \mathbf{x}_i , is calculated as:

$$\Omega_i = (\hat{\mathbf{x}} - \mathbf{x}_i)^T \mathbf{Q}_{\hat{\mathbf{x}}}^{-1} (\hat{\mathbf{x}} - \mathbf{x}_i) \quad (2.19)$$

Equation (2.19) represents the distance between the float vector and candidate set under investigation. The candidate set with the lowest square norm is considered the best set. To help discriminate between the best candidate set and the second best candidate, a ratio test is performed. The ratio test is defined as:

$$\text{ratio} = \frac{\Omega_{2^{nd} \text{ best}}}{\Omega_{\text{best}}} \quad (2.20)$$

Generally, if the ratio is consistently above a certain threshold, the ambiguity set corresponding with Ω_{best} is considered the true set. The threshold is typically set at 2. An overview of Kalman filtering follows next.

2.4 *Kalman Filters*

The Kalman filter is an optimal linear estimator based on linear stochastic system models driven by Gaussian noise and implemented in a recursive data processing algorithm [32, 33]. The filter has two main components — a propagation cycle based on the dynamics model of the system of interest and an update cycle based on the model of the measurement errors. It incorporates information of known statistical properties to provide the best estimate of the variables of interest. To accomplish this estimate, it uses the statistical information regarding the uncertainty in the dynamics model, measurement errors, and initial conditions. It is not necessary to have previous measurements to provide an optimal state estimate.

The filter has two stages of operation: the propagation stage and the update stage. During the propagation stage, the filter predicts a new state estimated based on the state estimate at a previous time and the imperfect dynamics model. The update stage takes new measurements and incorporates them into the filter to produce the optimal state estimate. There are three basic assumptions needed for the Kalman filter to operate effectively: the system dynamics are linear, all noise processes are white (i.e., not time-correlated), and all noise processes are jointly Gaussian.

Since the measurement variables in GPS systems are not linear, a specialized Kalman filter implementation is needed. This generalized implementation is called an extended Kalman filter (EKF). The EKF is a first-order approximation of an optimal non-linear filter. The same assumptions and stages apply to the EKF as with the normal Kalman filter. More complete information on EKFs can be found in [33].

2.4.1 State Equations. A first-order approximation of the system dynamics for a non-linear state equation can be modeled as [33]:

$$\dot{\mathbf{x}}(t) = \mathbf{F}(t) \mathbf{x}(t) + \mathbf{G}(t) \mathbf{w}(t) \tag{2.21}$$

where

$\mathbf{x}(t)$ = the n -dimensional system state vector
 $\mathbf{F}(t)$ = the n -by- n state dynamics matrix
 $\mathbf{G}(t)$ = the n -by- s noise distribution matrix
 $\mathbf{w}(t)$ = the s -dimensional dynamics driving noise vector

As mentioned earlier, all noise processes are assumed white and Gaussian. Based on this assumption, the dynamics driving noise vector, $\mathbf{w}(t)$, mean and covariance are defined by:

$$\begin{aligned}
 E \{ \mathbf{w}(t) \} &= \mathbf{0} \\
 E \{ \mathbf{w}(t) \mathbf{w}^T(t + \tau) \} &= \mathbf{Q}(t) \delta(\tau)
 \end{aligned} \tag{2.22}$$

where $\delta(\cdot)$ is the Dirac delta function

The equivalent discrete-time system dynamics equation is:

$$\mathbf{x}(t_{i+1}) = \mathbf{\Phi}(t_{i+1}, t_i) \mathbf{x}(t_i) + \mathbf{w}_d(t_i) \tag{2.23}$$

where

$\mathbf{x}(t_i)$ = the n -dimensional system state vector
 $\mathbf{\Phi}(t_{i+1}, t_i)$ = the n -by- n state transition matrix associated
 with $\mathbf{F}(t)$ in Equation (2.21)

The discrete-time white Gaussian dynamics driving noise vector, \mathbf{w}_d , is defined as:

$$\mathbf{w}_d(t_i) = \int_{t_i}^{t_{i+1}} \mathbf{\Phi}(t_{i+1}, \tau) \mathbf{G}(\tau) \mathbf{w}(\tau) d\tau \tag{2.24}$$

with statistics:

$$\begin{aligned}
 E \{ \mathbf{w}_d(t_i) \} &= \mathbf{0} \\
 E \{ \mathbf{w}_d(t_i) \mathbf{w}_d^T(t_i) \} &= \mathbf{Q}_d(t_i) = \int_{t_i}^{t_{i+1}} \mathbf{\Phi}(t_{i+1}, \tau) \mathbf{G}(\tau) \mathbf{Q}(\tau) \mathbf{G}^T(\tau) \mathbf{\Phi}^T(t_{i+1}, \tau) d\tau \\
 E \{ \mathbf{w}_d(t_i) \mathbf{w}_d^T(t_j) \} &= \mathbf{0}, \quad t_i \neq t_j
 \end{aligned} \tag{2.25}$$

2.4.2 State Transition Matrix. The state transition matrix, Φ , is used by the EKF in this research (since the dynamics model is linear) to propagate the state vector forward from time t_i to time t_{i+1} . The discrete-time state transition matrix is related to the \mathbf{F} matrix, in that it satisfies the differential equation and initial condition:

$$\frac{d}{dt} [\Phi(t, t_0)] = \mathbf{F}(t) \Phi(t, t_0) \quad (2.26)$$

$$\Phi(t_0, t_0) = \mathbf{I} \quad (2.27)$$

Assuming \mathbf{F} is time invariant, then Φ can be expressed as a matrix exponential:

$$\Phi(t_{i+1}, t_i) = \Phi(t_{i+1} - t_i) = e^{\mathbf{F}(\Delta t)} \quad (2.28)$$

where $\Delta t = t_{i+1} - t_i$.

2.4.3 Measurement Model. In sampled data form, the EKF measurement equation is specified by:

$$\mathbf{z}(t_i) = \mathbf{h}[\mathbf{x}(t_i), t_i] + \mathbf{v}(t_i) \quad (2.29)$$

where

$\mathbf{z}(t_i) =$ the m -dimensional measurement vector at time t_i

$\mathbf{h}[\mathbf{x}(t_i), t_i] =$ the m -dimensional measurement model vector

$\mathbf{v}(t_i) =$ the m -dimensional discrete-time noise vector

Once more, the discrete time noise vector, $\mathbf{v}(t_i)$, is assumed to be white Gaussian noise with mean and covariance:

$$\begin{aligned} E \{ \mathbf{v}(t_i) \} &= \mathbf{0} \\ E \{ \mathbf{v}(t_i) \mathbf{v}^T(t_j) \} &= \begin{cases} \mathbf{R}(t_i) & t_i = t_j \\ \mathbf{0} & t_i \neq t_j \end{cases} \end{aligned} \quad (2.30)$$

where

$\mathbf{R}(t_i) =$ the m -by- m measurement error covariance matrix at time t_i

Since the GPS measurements used in this thesis are non-linear, the state and measurement equations must be linearized for use in the EKF. For linearization purposes, we can define $\mathbf{H}[t_i; \hat{\mathbf{x}}(t_i^-)] \equiv \left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{d\mathbf{x}} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)}$ where $\hat{\mathbf{x}}(t_i^-)$ is the state estimate at time t_i prior to incorporating the measurement at that time. A detailed discussion of how the first-order linear approximations were formed, based on perturbations about a nominal state trajectory, is developed in [33]. The final steps in the EKF filter development are the state propagation and measurement incorporation.

2.4.4 State Estimates and Error Covariance Propagation. Propagation of the filter takes place between measurement epochs. The following equation is used to propagate the state estimates:

$$\hat{\mathbf{x}}(t_{i+1}^-) = \Phi(t_{i+1}, t_i) \tag{2.31}$$

and to propagate the error covariance:

$$\mathbf{P}(t_{i+1}^-) = \Phi(t_{i+1}, t_i) \mathbf{P}(t_i^+) \Phi^T(t_{i+1}, t_i) + \mathbf{Q}_d(t_i) \tag{2.32}$$

where $\mathbf{P}(t_{i+1}^-)$ is the conditional covariance of $\mathbf{x}(t_{i+1})$ before the measurement $\mathbf{z}(t_{i+1})$ is incorporated and $\hat{\mathbf{x}}(t_i^+)$ is the state estimate at time t_i after the measurement incorporation at that time. The last step in the EKF is measurement incorporation.

2.4.5 Measurement Incorporation. Measurement incorporation is done each time a new set of measurements is available to the filter. The new measurements are incorporated to update the filter's state estimate. The time the new measurements became available is designated t_i , and a superscript “-” and “+” is used to indicate the state estimate before and after measurement incorporation, respectively.

To accomplish the update the following equations are used [32]:

$$\mathbf{A}(t_i) = \mathbf{H}(t_i) \mathbf{P}(t_i^-) \mathbf{H}^T(t_i) + \mathbf{R}(t_i) \quad (2.33)$$

$$\mathbf{K}(t_i) = \mathbf{P}(t_i^-) \mathbf{H}^T(t_i) \mathbf{A}^{-1}(t_i) \quad (2.34)$$

$$\mathbf{r}(t_i) = \mathbf{z}(t_i) - \mathbf{h} [\mathbf{x}(t_i^-), t_i] \quad (2.35)$$

$$\hat{\mathbf{x}}(t_i^+) = \hat{\mathbf{x}}(t_i^-) + \mathbf{K}(t_i) \mathbf{r}(t_i) \quad (2.36)$$

$$\mathbf{P}(t_i^+) = \mathbf{P}(t_i^-) - \mathbf{K}(t_i) \mathbf{H}(t_i) \mathbf{P}(t_i^-) \quad (2.37)$$

When the filter is designed correctly and properly tuned, the residual vector, $\mathbf{r}(t_i)$, is zero-mean with a predicted covariance of $\mathbf{A}(t_i)$. The optimal state estimates, $\hat{\mathbf{x}}(t_i^+)$, and their associated filter computed covariance, $\mathbf{P}(t_i^+)$, are the primary outputs from the EKF. Now that the EKF has been described, the adaptations required for the MMAE are discussed.

2.5 Multiple Model Adaptive Estimation

All of the multiple model adaptive estimation theory used in this thesis was based on previous research done by Henderson [20]. The theory is repeated here for completeness and is similar to that shown in [20].

Most MMAEs are designed to look for changes within the system of interest. The MMAE used in this thesis is based on a multiple model filter algorithm for GPS carrier-phase ambiguity resolution derived in [33]. Figure 2.3 shows a block diagram of the multiple model filter algorithm.

There are a total of K individual filters in the MMAE. The individual filters are called elemental filters (EF). All K filters are given the same set of measurements at time t_i . However, each EF hypothesizes a different set of carrier-phase integer ambiguities as represented by the parameter vector, \mathbf{a}_k . After each update cycle, the residual vector, $\mathbf{r}_k(t_i)$, and its covariance matrix, $\mathbf{A}_k(t_i)$ indicate how well EF filter

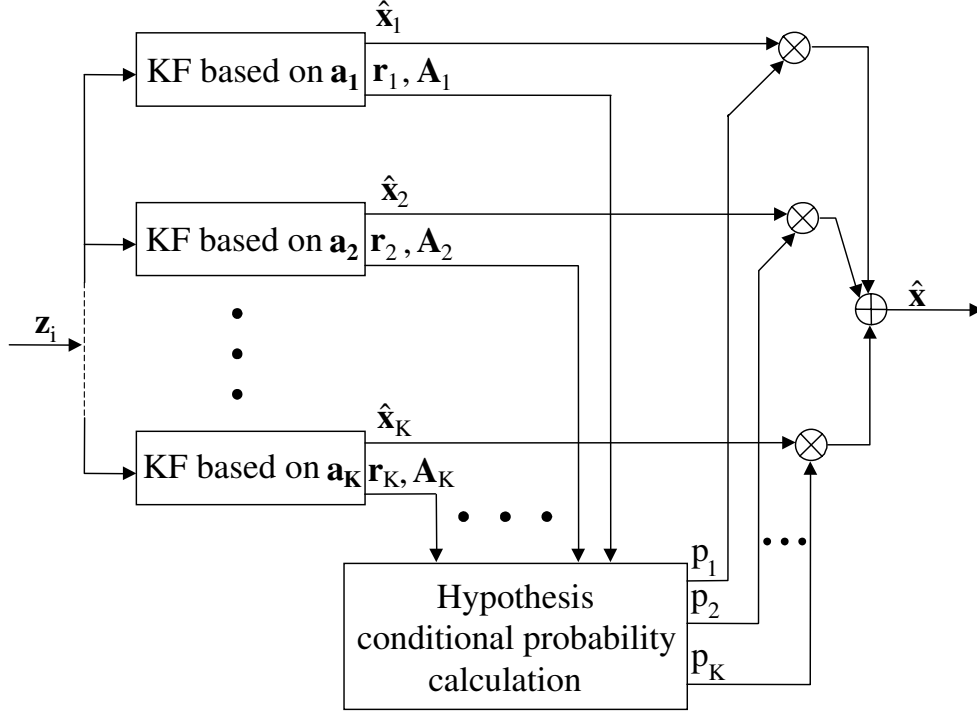


Figure 2.3: Multiple Model Filter Algorithm [20]

k models the true system. In general, the EF with the lowest residuals most closely matches the true system.

The EFs are propagated and updated in the same fashion as described in Sections 2.4.4 and 2.4.5. A k subscript is used on each vector and matrix to denote the particular elemental filter of interest. At each epoch, all K filters generate an optimal state estimate, $\hat{\mathbf{x}}_k(t_i^+)$, covariance, $\mathbf{P}_k(t_i^+)$, residual vector, $\mathbf{r}_k(t_i)$, and covariance, $\mathbf{A}_k(t_i)$, based on its own parameter vector. The residual vector, $\mathbf{r}_k(t_i)$, and its covariance, $\mathbf{A}_k(t_i)$, are of particular use in the MMAE, especially the conditional probability calculations.

2.5.1 Conditional Probability Calculation. As seen in Figure 2.3, all K filters have a conditional probability calculation performed at each epoch. The conditional probability calculation at time, t_i , is specified by:

$$p_k(t_i) = \text{Prob}(\mathbf{a} = \mathbf{a}_k | \mathbf{Z}(t_i) = \mathbf{Z}_i) \quad (2.38)$$

The conditional probability is the probability that the random variable, \mathbf{a} , representing the system characteristics is the realization of the variable in the k^{th} elemental filter, \mathbf{a}_k , given the time history of measurements up to and including that taken at time t_i . The time history of measurements is represented as:

$$\mathbf{Z}(t_i) = \left[\mathbf{z}^T(t_1) : \mathbf{z}^T(t_2) : \dots : \mathbf{z}^T(t_i) \right]^T \quad (2.39)$$

where $\mathbf{Z}(t_i)$ is the vector of actually observed realizations of that history of measurements.

From [33], the conditional probability is calculated using:

$$p_k(t_i) = \frac{f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i | \mathbf{a}_k, \mathbf{Z}_{i-1}) p_k(t_{i-1})}{\sum_{j=1}^K f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i | \mathbf{a}_j, \mathbf{Z}_{i-1}) p_j(t_{i-1})} \quad (2.40)$$

The denominator in Equation (2.40) acts as a scaling function to ensure the probability of each EF is properly defined in the sense that:

$$p_k(t_i) \geq 0 \text{ for all } k \text{ and } \sum_{j=1}^K p_j(t_i) = 1 \quad (2.41)$$

The conditional density function (the first part of the numerator) in Equation (2.40) is calculated by:

$$\begin{aligned} f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i | \mathbf{a}_k, \mathbf{Z}_{i-1}) &= \frac{1}{(2\pi)^{m/2} \sqrt{|\mathbf{A}_k(t_i)|}} e^{\{\bullet\}} \\ \{\bullet\} &= -\frac{1}{2} \mathbf{r}_k^T(t_i) \mathbf{A}_k^{-1}(t_i) \mathbf{r}_k(t_i) \end{aligned} \quad (2.42)$$

Equation (2.42) represents the conditional density function for the measurements at time t_i of the k^{th} EF, conditioned on \mathbf{a}_k and the time history of measurements up to time t_{i-1} , where $\mathbf{r}_k(t_i)$ and $\mathbf{A}_k(t_i)$ are given by Equations (2.35) and (2.33). The m in Equation (2.42) is the length of the measurement vector. If designed correctly, the EF representing the true system will have zero-mean white Gaussian residuals with consistently smaller values than the other EFs. Therefore, based on Equation (2.40),

the probability of the correct filter increases while the probability of the other filters decreases.

2.5.1.1 Lower Probability Bounds. Based on the recursive nature of the conditional probability calculation, one change to the probability calculation is needed. From Equation (2.40), the EF's conditional probability at time t_i is based on the current epoch's residuals and covariance and every previous probability calculated for EF k up to and including time t_{i-1} . Thus, if an EF is assigned a probability of zero at any time, the EF's conditional probability will remain at zero from that time forward due to the iteration form of Equation (2.40). To alleviate this problem, a lower probability bound, ε , is established for each filter. Normally, ε is between 0.001 and 0.01 [20]. If the calculated EF's conditional probability is less than the lower bound, the EF is assigned a probability of ε and all probabilities are then rescaled to add to one. Selecting the appropriate value for ε is important for accurate conditional probability calculations. If the value of the lower bound is too large, it will inappropriately weight the wrong filter. On the other hand, very small values of ε slow down the response of the MMAE to true parameter or system changes [33].

2.5.2 Overall MMAE State Estimate. The output of the MMAE algorithm is a probabilistically weighted average of the outputs from the K EFs. The overall state estimate of the MMAE is calculated as:

$$\hat{\mathbf{x}}(t_i^+) = \sum_{j=1}^K \hat{\mathbf{x}}_j(t_i^+) p_j(t_i) \quad (2.43)$$

and the conditional covariance of the MMAE is calculated as:

$$\mathbf{P}(t_i^+) = \sum_{j=1}^K p_j(t_i) \left\{ \mathbf{P}_j(t_i^+) + [\hat{\mathbf{x}}_j(t_i^+) - \hat{\mathbf{x}}(t_i^+)] [\hat{\mathbf{x}}_j(t_i^+) - \hat{\mathbf{x}}(t_i^+)]^T \right\} \quad (2.44)$$

Note that the output of the MMAE includes the outputs from all of the EFs, even those that are wrong. Of course, the incorrect filters should have a very low

probability, while the correct filter a very high probability. If that is the case, then the MMAE blended solution is still relatively accurate based primarily on the EF that matches the true parameter \mathbf{a} (the correct ambiguity set, in this case).

This concludes the discussion of the MMAE. More detailed information on MMAEs can be found in [20, 32, 33, 44].

2.6 Summary

Chapter II has described the theory required to form a basis for the remaining chapters. First, an overview of GPS was given with emphasis on the GPS signal structure and different measurement observables. After GPS, the LAMBDA method for generating candidate ambiguity sets for the carrier-phase measurement observables was discussed. Next, Kalman filtering was discussed, including the extended Kalman filter (EKF) and multiple model adaptive estimation (MMAE). The implementation of the theory presented in this chapter will be put into practice in Chapter III.

III. Hardware Design and Algorithm Development

3.1 Overview

This chapter describes the hardware configuration, overall algorithm, system models, and techniques used throughout this thesis. First, an overall system description is given, including the hardware, followed by a more detailed discussion of each sub-system. Next, the Kalman filter states, equations, and other new techniques developed are given. Finally, areas of additional functionality are explained. This chapter closely follows [20] in both form and content.

3.2 System Configuration

The system was designed to provide a centimeter-level relative position between two aircraft for use in applications such as automated aerial-refueling or close formation flying. A real-time kinematic (RTK) precise differential Global Positioning Satellite (DGPS) algorithm was developed to fulfill the requirement. Design and initial testing of the system was conducted in the navigation laboratory at the Air Force Institute of Technology (AFIT) including both static ground tests and dynamic driving tests. Flight tests occurred at Edwards Air Force Base (EAFB), CA.

Four test units were built. The hardware consisted of a PC-104 stack with 256 MB of memory and a 600 MHz processor. Included on the stack, was a PC-104 version of the Freewave™ modem built by MicroBee Systems, Incorporated and a Javad® JNS100 GPS receiver. The PC-104 stack was housed within a Diamond System Pandora Enclosure measuring 5" long, 5.5" wide, and 5.75" high. Lastly, a special power supply was used which would work with the 28V aircraft electrical system and absorb any aircraft power transients. Specifications for the Javad® JNS100, Freewave™ modem, and the Diamond System enclosure can be found in [11, 35, 39] respectively.

The test units had all of the functionality of a laptop computer. This included four serial ports, mouse and keyboard ports, monitor hook-up, two USB ports, and Ethernet. Also, for use on aircraft, a 1553 bus connector was included. In the lead

aircraft's test unit, an MT9 inertial navigation unit (INU) was attached inside the test unit on the back plate of the aluminum hard shell via velcro [38]. The velcro was used for easy installation and removal, but had the added benefit of dampening high frequency vibrations. There was some concern about high frequency vibrations from the engines on the C-12C. The hardware rack in the C-12C had rubber buffers to help prevent any vibrations from propagating to the top platform, but the velcro was additional protection. Figure 3.1 shows the test unit installed on the test aircraft, a C-12C. Figure 3.2 depicts the inside back-plate of the lead test unit with the MT9 installed. Two test units were sent to EAFB for flight testing and two were kept at AFIT for development and ground testing.



Figure 3.1: Test Unit Installed on C-12



Figure 3.2: Inside Back-plate of Test Unit

The Javad® receiver had numerous configuration options. Table 3.1 shows some of the receiver options used during testing. The multipath option significantly helped

Table 3.1: Javad® GPS Receiver Configuration

Option	Value
Multipath Reduction	On
Loop Bandwidth	25
Loop Order	3
WAAS	On

reduce multipath errors on the code measurements, especially in multipath prone environments such as near tall buildings. Javad® claimed a reduction of up to 50 percent on the multipath errors was possible with the multipath reduction option turned on.

During lab testing, it was discovered that the lead Freewave™ modem would retransmit packets which were not acknowledged by the wing modem. There were

occasions when the wing modem would receive a packet, but the acknowledgment was “lost”. Once the lead modem began retransmitting packets, the system delay became unacceptably large. Since the algorithm ran in real-time, if a packet was lost, there was no need to resend old information. Therefore, the retransmit feature in the lead modem was turned off. If a packet was lost, the relative position update skipped an epoch until two synchronized messages were available. Each modem was set at a 56K baud rate, the highest available. The overall system configuration and message flow can be seen in Figure 3.3.

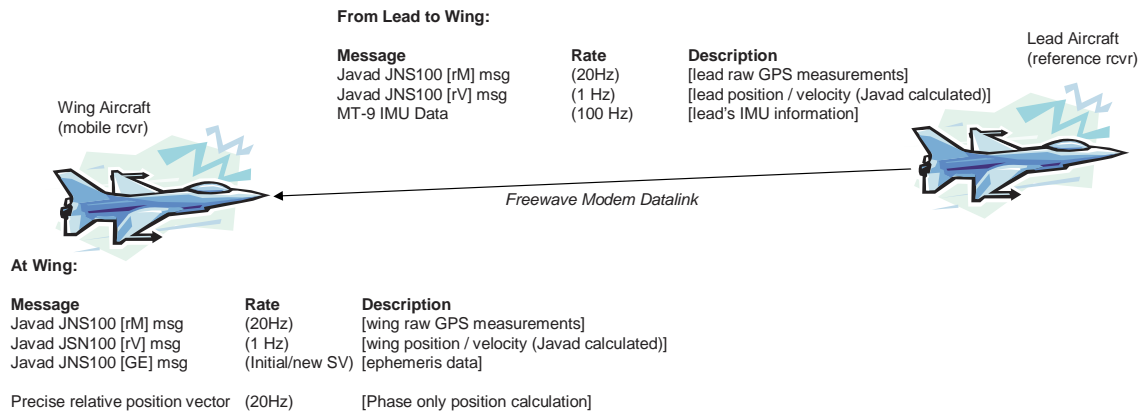


Figure 3.3: Overall Message Flow

While in the navigation lab, the test units had dedicated monitors, mice and keyboards. During the driving and flight testing, a secure shell was used to login to the test units via a laptop and ethernet cable. For the driving tests, two of the author’s personal vehicles were used. Two C-12Cs were used for the initial flight tests in April 2005. The second set of flight tests, in October 2005, used one of the previously modified C-12C and a highly modified Calspan Learjet 24 as the two test aircraft. The C-12C aircraft had previously been modified for other research. The modifications included a GPS Aided Inertial Navigation Reference (GAINR) [34] GPS receiver for use as the truth source. However, the GAINR truth source was only accurate to a couple of feet. Therefore, a commercial product was used to determine the ambiguities and relative positions during post processing. GrafMov, developed

by Waypoint Consulting Incorporated, was chosen for its ability to determine the double-difference ambiguities with a roving base station [36].

3.2.1 Data Collection. Numerous output files were generated by the code during a data run. First, the parsed and raw Javad[®] messages were saved for post-processing capability. Also, a file containing a GPS time stamp with the associated system delay, relative position and MT9 information was saved. Finally, three filter-specific files were created that contained information on the floating filter, the sum-of-squared (SOS) residuals for every candidate set, and a time-stamped list of minimum indicator ambiguity sets chosen throughout the data run. Most of the data files were stored in binary format to reduce file size. All of the binary files were decoded using MATLAB[®] for easier processing and plotting [37].

3.3 Overall Algorithm Structure

Figure 3.4 gives a visual representation of the overall algorithmic flow. There are four main areas of interest: the floating point DGPS Kalman filter, the candidate ambiguity set generation, the Multiple Model Adaptive Estimation (MMAE) design, and the minimum indicator design. All of these areas are described in the following sections, along with some other areas of interest not specifically shown on the Figure 3.4 graphic.

3.4 Floating Point DGPS Kalman Filter

The first section of interest is the floating point DGPS Kalman filter. Since there was only one floating point DGPS Kalman filter in the software, it will be referred to as the floating point filter from here forward. There were three primary functions of the floating point filter. First, it estimated the relative position, and optionally the relative velocity and relative acceleration between the phase center of the two GPS antennas. Second, it fed the LAMBDA algorithm an estimate of the double-difference phase ambiguities and the associated covariance matrix. Finally, the filter was used

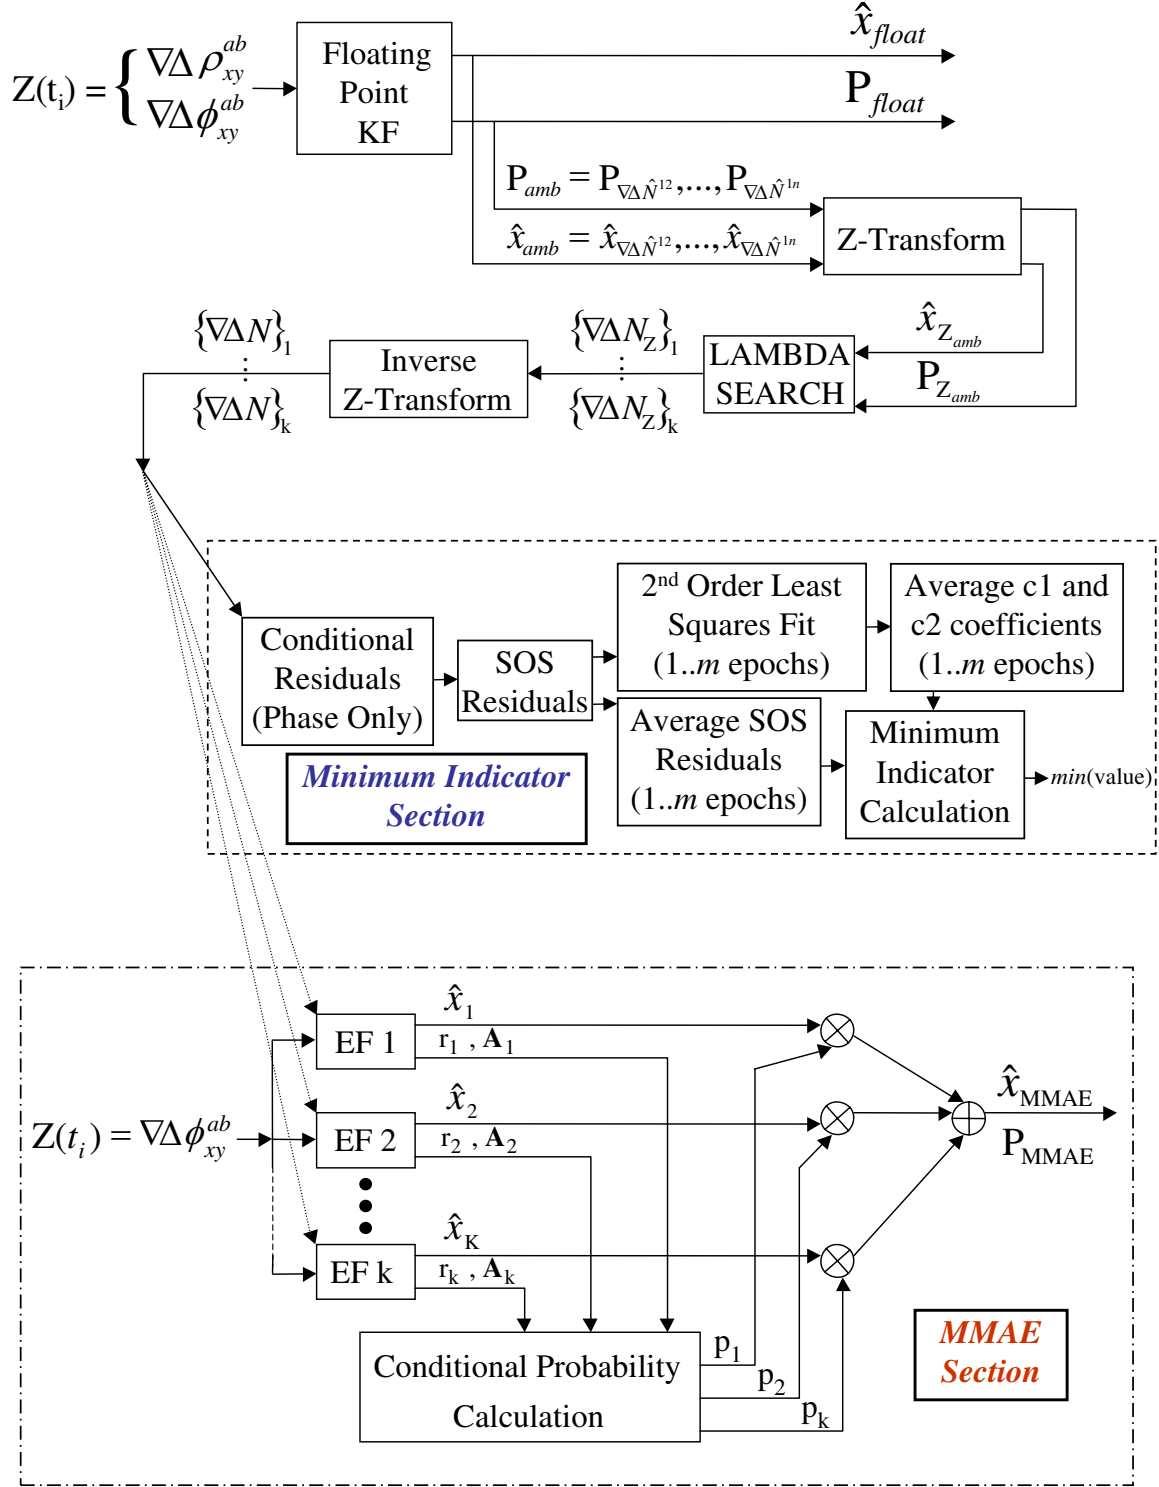


Figure 3.4: Overall Algorithm Flow

to initialize, or reinitialize, the elemental filters in the MMAE. The structure of the floating point filter mimicked the construction found in [47], with minor modifications for a roving reference station. For the purposes of this thesis, the mobile receiver was the wing aircraft and the reference receiver was the lead aircraft.

The filter was originally coded in MATLAB®. However, due to MATLAB®'s relatively slow processing speed compared with C++, and the need to use Linux as the operating system, the filter was coded using C++ and KDevelop [28] in SUSE LINUX Professional 9.2 [42]. The driving factor for using Linux was the operating system's ability to access the serial ports in a reliable fashion. Under the original MATLAB® and Windows® environment, system maintenance tasks periodically overrode the serial I/O. With a filter running at 20 Hz, this was unacceptable. Joffrion had similar problems using MATLAB® and Windows® in his thesis [22].

Since the algorithm was designed to run in real-time, it was necessary to make the floating point filter a background thread during execution. This allowed the program to maintain the 20 Hz precise DGPS relative position output (which was not based on the floating filter) without waiting for the floating point filter to finish its calculations at each one second interval. The floating point filter ran at a one Hz rate. Using more than a one Hz update rate for the floating point filter broke the Extended Kalman Filter (EKF) assumption that the errors in the measurements were white and Gaussian in nature. Errors such as multipath are time-correlated, especially at short intervals. A detailed view of the message flow can be seen in Figure 3.5.

The Javad® JNS100 receivers used in this thesis were single-frequency, L1 only. Originally, the Javad® receivers were chosen for their ability to output raw GPS measurements at 100 Hz rate. The program was designed to run at this rate, but the Freewave™ modem data link would not reliably transmit all message traffic from the lead to the wing unit during ground testing. The 20 Hz rate was chosen for its reliability and still relatively high output rate. The single-frequency receiver made determining the double-difference ambiguities more difficult than if it were a dual-

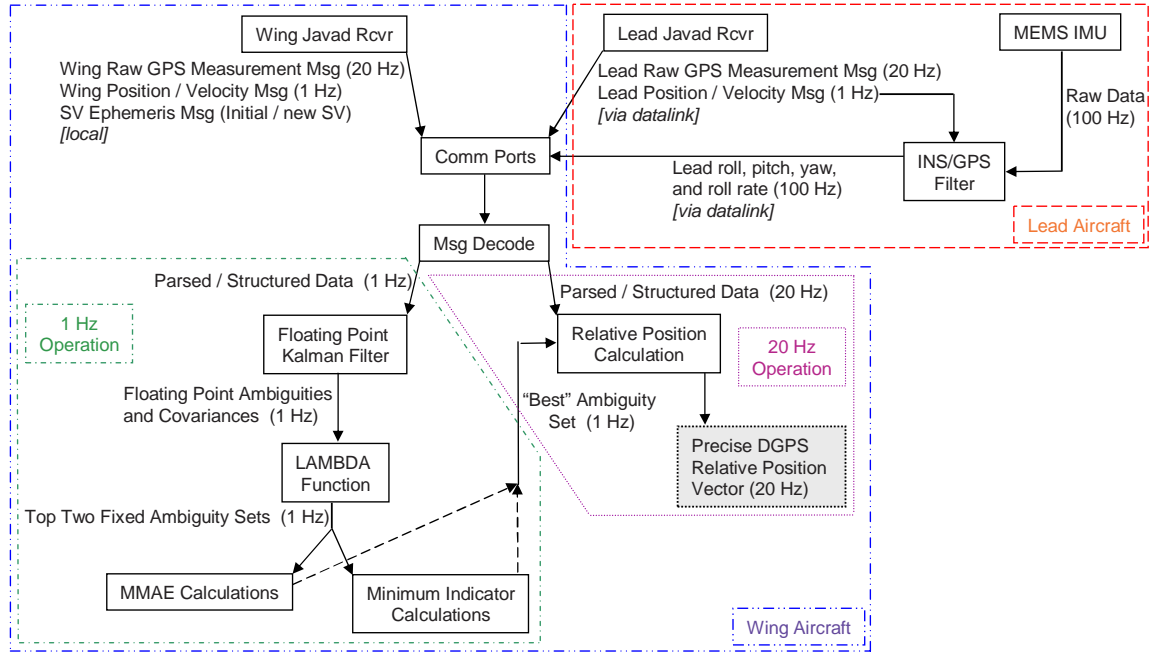


Figure 3.5: Algorithm Message Flow

frequency receiver. For this reason, a two-minute convergence time was given for the floating point filter before any floating point ambiguities were sent to the LAMBDA function. The two-minute wait time was determined empirically. By waiting two minutes and giving the floating point filter time to converge, the number of erroneous ambiguity sets was significantly reduced. Usually, the correct ambiguity set was within the top ten sets returned from the LAMBDA function. By reducing the number of incorrect fixed ambiguity sets, the computational load was reduced and the MMAE weighted position more accurate.

Another critical aspect of the floating point filter was timing. The filter needed raw measurements from both receivers corresponding to the same GPS time. Timing was performed by using a roving bank of ten measurements from each receiver. A current position marker and a synchronization marker was used for each list. The current position marker pointed to the most current measurement in the list, while the synchronization marker pointed to the latest time that appeared in both the lead's and wing's measurement list. Since the wing aircraft housed the floating point filter and

the lead aircraft messages had to go through multiple serial ports before arriving at the wing aircraft, the wing aircraft current position marker would usually be one message ahead of the lead aircraft. On average, the total delay from when a GPS message arrived until the precise relative position was calculated, was 80 milliseconds. This value took into account all overhead, including parsing, synchronizing, and calculation time.

The time delay was calculated using the computer clock. When the program started, a running timer began as a reference point. Subsequently, when a Javad® position message, rV message, was received on the wing aircraft, the system receive time was noted. After all of the required processing was completed, and the precise relative position message was copied to the output stream, the system sent time was noted. The total delay was the difference between the system sent time and the system receive time.

Finally, it should be noted that all computationally intensive functions such as the floating point filter, LAMBDA function, minimum indicator calculation and MMAE calculations were calculated at a one Hz rate in a separate thread from the main program. However, critical functions such as parsing incoming messages, organizing data and precise DGPS relative positioning were accomplished at a 20 Hz rate. More detailed discussions of each area will be addressed later in the chapter.

3.4.1 Floating Point Filter Model Equations. The floating point DGPS filter varies, depending upon user input. To facilitate quicker design and testing, an initialization file was created. Based on the initialization file, the floating point filter could either have $9 + (n - 1)$ states with the relative velocity and acceleration states included or $3 + (n - 1)$ states with relative position only; n was the number of visible satellites (SV). Also, since the raw GPS measurements from the Javad® JNS100 were in the Earth Centered Earth Fixed (ECEF) frame, all filter calculations were conducted in the ECEF frame. The overall state vector for the floating point

filter when the relative velocity and acceleration states were included was defined as:

$$\mathbf{x} = \left[\Delta X \ \Delta Y \ \Delta Z \ \Delta \dot{X} \ \Delta \dot{Y} \ \Delta \dot{Z} \ \Delta \ddot{X} \ \Delta \ddot{Y} \ \Delta \ddot{Z} \ \Delta \nabla N^{12} \dots \Delta \nabla N^{1n} \right]^T \quad (3.1)$$

where

$$\begin{aligned} x_1 &= \Delta X = \text{Relative ECEF X position (m)} \\ x_2 &= \Delta Y = \text{Relative ECEF Y position (m)} \\ x_3 &= \Delta Z = \text{Relative ECEF Z position (m)} \\ x_4 &= \Delta \dot{X} = \text{Relative ECEF X velocity (m/s)} \\ x_5 &= \Delta \dot{Y} = \text{Relative ECEF Y velocity (m/s)} \\ x_6 &= \Delta \dot{Z} = \text{Relative ECEF Z velocity (m/s)} \\ x_7 &= \Delta \ddot{X} = \text{Relative ECEF X acceleration (m/s}^2\text{)} \\ x_8 &= \Delta \ddot{Y} = \text{Relative ECEF Y acceleration (m/s}^2\text{)} \\ x_9 &= \Delta \ddot{Z} = \text{Relative ECEF Z acceleration (m/s}^2\text{)} \\ x_{10} &= \Delta \nabla N^{12} = \text{double-difference ambiguity} \\ &\quad \text{between satellites 1 \& 2 (cycles)} \\ &\quad \vdots \\ x_{9+(n-1)} &= \Delta \nabla N^{1n} = \text{double-difference ambiguity} \\ &\quad \text{between satellites 1 \& n (cycles)} \end{aligned}$$

When the floating filter was running in position only mode, there were no velocity and acceleration states and the double-difference ambiguity states started in x_4 .

During testing, the filter generally ran with the relative velocity and acceleration states turned on. Based on a First Order Gauss Markov Acceleration (FOGMA) model, the relative position and velocity states' dynamics were defined as:

$$\begin{aligned} \dot{x}_1 &= \Delta \dot{X} = x_4 & \dot{x}_2 &= \Delta \dot{Y} = x_5 & \dot{x}_3 &= \Delta \dot{Z} = x_6 \\ \dot{x}_4 &= \Delta \ddot{X} = x_7 & \dot{x}_5 &= \Delta \ddot{Y} = x_8 & \dot{x}_6 &= \Delta \ddot{Z} = x_9 \end{aligned} \quad (3.2)$$

If the relative velocity and acceleration states were off, x_4 thru x_9 changed to the ambiguity state estimates (assuming there were at least six ambiguity states).

The relative acceleration states' dynamics were modeled as First Order Gauss Markov processes [47] as follows:

$$\begin{aligned}\dot{x}_7 &= (-1/T_a) x_7 + w_{a_1}(t) \\ \dot{x}_8 &= (-1/T_a) x_8 + w_{a_2}(t) \\ \dot{x}_9 &= (-1/T_a) x_9 + w_{a_3}(t)\end{aligned}\tag{3.3}$$

The acceleration state process noises were defined as:

$$E \{w_{a_1}(t)w_{a_1}(t + \tau)\} = E \{w_{a_2}(t)w_{a_2}(t + \tau)\} = E \{w_{a_3}(t)w_{a_3}(t + \tau)\} = \frac{2\sigma_a^2}{T_a}\delta(\tau) = q_a\delta(\tau) \tag{3.4}$$

The time constant, T_a , and the variance, σ_a^2 , were user defined tuning values. The initialization file allowed for either static or dynamic setting of these values. In the case of static setting, the user input the values for both of these variables and they did not change throughout the data run. Typical static values for T_a and σ_a^2 are shown in Table 3.3. In dynamic mode, the T_a and σ_a^2 changed based on the total acceleration, $\sqrt{x_7^2 + x_8^2 + x_9^2}$, after measurement incorporation. As the total acceleration increases, T_a generally decreases, since the higher accelerations generally don't last long. In the same manner σ_a^2 generally increases since it is harder for the FOGMA model to accurately determine the true accelerations during higher maneuvering periods. The above generalizations are especially true of the aircraft used during testing in this thesis.

When dynamic T_a and σ_a^2 mode was commanded via the initialization file, the time constant and acceleration variance followed the progression in Table 3.2, based on the magnitude of the total acceleration (TA) vector.

Table 3.2: Dynamic T_a and σ_a^2 Progression

Total Acceleration (m/sec ²)	T_a (sec)	σ_a^2 (m ² /sec ⁴)	Resultant q_a (m ² /sec ⁵)
$TA \leq 1$	1000	25	0.05
$1 < TA \leq 2$	500	49	0.196
$2 < TA \leq 4$	20	49	4.9
$4 < TA \leq 6$	5	100	40
$6 < TA \leq 8$	2	169	169
$TA > 8$	1	169	338

The use of dynamic T_a and σ_a^2 values was restricted to post-processing only. The values set forth in Table 3.2 have not been flight tested, and are only initial estimates of potentially good

candidate values. They could change depending upon the application and vehicle dynamics. More testing is needed to validate other potential total acceleration break points, and the T_a and σ_a^2 values.

Unlike the model set forth in [47], there were no clock error states. Since the floating filter used double-differenced pseudoranges, the user clock errors cancel and no clock estimates were needed. The filter also used double-differenced carrier-phase measurements. In the software, the phase measurements could be L1, L2, or widelane. However, since the Javad[®] JNS100 receivers were only L1 capable, the bulk of the results in this thesis use L1 data. (An exception was one ground test, described in Chapter IV, which shows the potential benefits of using widelane measurement observables, namely quicker resolution of the carrier-phase ambiguities.) The ability for the filter to use widelane measurements was also an option in the initialization file. As mentioned before, in order to have precise relative positioning between the two receivers, the carrier-phase ambiguities must be resolved. The final states in the filter were used to estimate these ambiguities. By definition, the ambiguities were constant biases in the phase measurements (assuming no loss of lock or cycle slips). However, if they were modeled as constant biases, the filter may have converged to the wrong ambiguity set and would not be able to change as measurements dictated. To prevent this, the ambiguity states were modeled as random walks (with very small process noise). The carrier-phase ambiguity states dynamics were defined as:

$$\dot{x}_{10} = w_{\Delta\nabla N^{12}} \quad \dot{x}_{11} = w_{\Delta\nabla N^{13}} \quad \cdots \quad \dot{x}_{9+(n-1)} = w_{\Delta\nabla N^{1n}} \quad (3.5)$$

and the process noise defined as:

$$E \{w_{\Delta\nabla N^{12}}(t)w_{\Delta\nabla N^{12}}(t + \tau)\} = \cdots = E \{w_{\Delta\nabla N^{1n}}(t)w_{\Delta\nabla N^{1n}}(t + \tau)\} = q_N\delta(\tau) \quad (3.6)$$

The linear dynamics equation for the floating point filter was defined by:

$$\dot{\mathbf{x}}(t) = \mathbf{F}(t)\mathbf{x}(t) + \mathbf{G}(t)\mathbf{w}(t) \quad (3.7)$$

and represented in matrix form with:

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \\ \dot{x}_4 \\ \dot{x}_5 \\ \dot{x}_6 \\ \dot{x}_7 \\ \dot{x}_8 \\ \dot{x}_9 \\ \dot{x}_{10} \\ \dot{x}_{11} \\ \vdots \\ \dot{x}_{9+(n-1)} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \\ x_9 \\ x_{10} \\ x_{11} \\ \vdots \\ x_{9+(n-1)} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ w_{a_1} \\ w_{a_2} \\ w_{a_3} \\ w_{\Delta \nabla N^{12}} \\ w_{\Delta \nabla N^{13}} \\ \vdots \\ w_{\Delta \nabla N^{1n}} \end{bmatrix} \quad (3.8)$$

The T_a in Equation (3.8) was the time constant described earlier for the FOGMA model. Also, $\mathbf{G}(t)$, the Kalman gain matrix for the process noise was an identity matrix. When the floating point filter was in position-only mode, the dynamics matrix, $\mathbf{F}(t)$, was all zeros. Therefore, the system was driven entirely by process noise. The other difference in the position-only mode was the addition of process noise on the three position states. In this case, there would be w_{p_1} , w_{p_2} and w_{p_3} , respectively, in the first three positions of the white noise, $\mathbf{w}(t)$, vector.

Using the $\mathbf{w}(t)$ vector and assuming $E\{\mathbf{w}\} = \mathbf{0}$, a dynamics noise matrix, \mathbf{Q} , was defined as:

$$\mathbf{Q}(t) \delta(\tau) = E\{\mathbf{w}(t) \mathbf{w}^T(t + \tau)\} \quad (3.9)$$

in matrix form:

$$\mathbf{Q} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & q_a & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & q_a & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q_a & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q_N & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & q_N & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & q_N \end{bmatrix} \quad (3.10)$$

If the floating filter was in position-only mode, rows and columns four thru nine were removed and a q_p term went in $\mathbf{Q}_{1,1}$, $\mathbf{Q}_{2,2}$ and $\mathbf{Q}_{3,3}$. The remaining rows with q_N values remained the same.

Table 3.3 displays the common noise values used throughout the thesis for both ground and air testing. However, the time constant, T_a , and variance, σ_a^2 , for the acceleration states could dynamically change if the initialization file was set accordingly. Finally, the position noise, q_p , and acceleration values were mutually exclusive. The filter either had the acceleration states, in which case all of the acceleration values applied, or it was in position-only mode and only the position noise applied. In either case, the ambiguity noise remained the same.

Table 3.3: Static Floating Point Noise Values

Term	Definition	Value
σ_a^2	Acceleration mean squared value	$(15 \text{ m/sec}^2)^2$
T_a	Acceleration time constant	3 sec
q_a	Acceleration noise	$150 \text{ m}^2/\text{sec}^5$
q_p	Position noise	$200 \text{ m}^2/\text{sec}$
q_N	Ambiguity noise	$1.1 \times 10^{-4} \text{ cycles}^2/\text{sec}$

For the floating point filter to function efficiently and converge to the correct relative position and double-differenced phase ambiguities, approximate initial conditions were needed in the state vector. The relative position states (x_1, x_2, x_3) were initialized using position messages from the JNS100

receivers (rV message). The actual relative position calculation was defined as wing minus lead. The Javad® receivers were Wide Area Augmentation System (WAAS) capable. Therefore, their single-point positioning algorithms were usually within 1-2 meters, which was sufficiently adequate for the Kalman filter initialization. The double-differenced phase ambiguity states $(x_{10}, x_{11}, x_{9+(n-1)})$ were initialized using a code-carrier difference described by the following equation:

$$x_{1n} = \Delta\nabla N^{1n} \approx \Delta\nabla\phi^{1n} - \frac{\Delta\nabla\rho^{1n}}{\lambda} \quad (3.11)$$

where

$$\begin{aligned} \Delta\nabla\phi^{1n} &= \text{double-differenced phase between} \\ &\quad \text{satellites 1 \& n (cycles, L1 or widelane)} \\ \Delta\nabla\rho^{1n} &= \text{double-differenced pseudorange between satellites 1 \& n (m)} \\ \lambda &= \text{carrier wavelength (m/cycle, L1 or widelane)} \end{aligned}$$

All other states were initialized to zero making the initial state vector:

$$\mathbf{x}(t_0) = \left[\Delta X_0 \quad \Delta Y_0 \quad \Delta Z_0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad \Delta\nabla N_0^{12} \quad \Delta\nabla N_0^{13} \quad \dots \quad \Delta\nabla N_0^{1n} \right]^T \quad (3.12)$$

The final part of the initial floating filter development was the covariance matrix, \mathbf{P} . Since accurate initial cross-covariance values were not known, all cross-covariance terms were set to zero. As the filter converged, it would populate these values while it “learned” how each of the states were related to the others. The initial \mathbf{P} matrix was given by:

$$\mathbf{P}(t_0) = \begin{bmatrix} \sigma_{\Delta X}^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & \sigma_{\Delta Y}^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & \sigma_{\Delta Z}^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & \sigma_{\Delta \dot{X}}^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & \sigma_{\Delta \dot{Y}}^2 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_{\Delta \dot{Z}}^2 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\Delta \ddot{X}}^2 & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\Delta \ddot{Y}}^2 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\Delta \ddot{Z}}^2 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\Delta \nabla N^{12}}^2 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \sigma_{\Delta \nabla N^{13}}^2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & \sigma_{\Delta \nabla N^{1n}}^2 \end{bmatrix} \quad (3.13)$$

The initial values for the diagonal of the $\mathbf{P}(t_0)$ matrix are specified in Table 3.4 below.

Table 3.4: Initial Floating Point Filter Covariance Values

Term	Definition	Value
$\sigma_{\Delta X, \Delta Y, \Delta Z}^2$	Position state variance	$(5.0 \text{ m})^2$
$\sigma_{\Delta \dot{X}, \Delta \dot{Y}, \Delta \dot{Z}}^2$	Velocity state variance	$(10.0 \text{ m/sec})^2$
$\sigma_{\Delta \ddot{X}, \Delta \ddot{Y}, \Delta \ddot{Z}}^2$	Acceleration state variance	$(5.0 \text{ m/sec}^2)^2$
$\sigma_{\Delta \nabla N^{12}}^2 = \cdots \sigma_{\Delta \nabla N^{1n}}^2$	Ambiguity state variance	$(\frac{50}{\lambda} \text{ cycles})^2$

3.4.2 Differential GPS Measurement Model. Since the GPS measurements used by the floating filter were non-linear, an extended Kalman filter formulation was chosen with non-linear measurements. The non-linear measurement equation was specified by:

$$\mathbf{z}(t_i) = \mathbf{h}[\mathbf{x}(t_i), t_i] + \mathbf{v}(t_i) \quad (3.14)$$

where

$\mathbf{z}(t_i)$ = measurement vector at time t_i

$\mathbf{v}(t_i)$ = Zero-mean white Gaussian vector process of covariance \mathbf{R}

with the assumption:

$$E \{ \mathbf{v}(t_i) \mathbf{v}^T(t_j) \} = \begin{cases} \mathbf{R} & t_i = t_j \\ \mathbf{0} & t_i \neq t_j \end{cases} \quad (3.15)$$

Based on using both double-differenced pseudoranges and double-differenced phases, the measurement vector, $\mathbf{z}(t_i)$, had $2n - 2$ states (where n was the number of measurements):

$$\mathbf{z}(t_i) = \begin{bmatrix} \Delta\nabla\rho_{xy}^{12} & \Delta\nabla\rho_{xy}^{13} & \cdots & \Delta\nabla\rho_{xy}^{1n} & \Delta\nabla\phi_{xy}^{12} & \Delta\nabla\phi_{xy}^{13} & \cdots & \Delta\nabla\phi_{xy}^{1n} \end{bmatrix}^T \quad (3.16)$$

where

$$\begin{aligned} \Delta\nabla\rho_{xy}^{12} &= \text{double-differenced code measurement between satellites 1 \& 2} \\ \Delta\nabla\rho_{xy}^{13} &= \text{double-differenced code measurement between satellites 1 \& 3} \\ &\vdots \\ \Delta\nabla\rho_{xy}^{1n} &= \text{double-differenced code measurement between satellites 1 \& } n \\ \Delta\nabla\phi_{xy}^{12} &= \text{double-differenced phase measurement between satellites 1 \& 2} \\ \Delta\nabla\phi_{xy}^{13} &= \text{double-differenced phase measurement between satellites 1 \& 3} \\ &\vdots \\ \Delta\nabla\phi_{xy}^{1n} &= \text{double-differenced phase measurement between satellites 1 \& } n \end{aligned}$$

From Section 2.2.2.1, the double-differenced code measurement was defined as:

$$\Delta\nabla\rho_{xy}^{ab} = \Delta\nabla r_{xy}^{ab} + \Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab} \quad (3.17)$$

By combining like terms the equation reduced to:

$$\Delta\nabla\rho_{xy}^{ab} = (r_y^a - r_x^a) - (r_y^b - r_x^b) + \nu_{\Delta\nabla\rho} \quad (3.18)$$

where

$$\begin{aligned} (r_y^a - r_x^a) - (r_y^b - r_x^b) &= \Delta\nabla r_{xy}^{ab} \\ \nu_{\Delta\nabla\rho} &= \Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab} \end{aligned}$$

The ν vector was assumed zero-mean, white and Gaussian in nature. The double-difference tropospheric and ionospheric errors were assumed to be zero since the baseline distance was so short (less than 5 km, typically less than 100 m). However, all the variables in the $\nu_{\Delta\nabla\rho}$ term did have some time correlation, except the noise variable, $\Delta\nabla v_{xy}^{ab}$. By using a sufficiently large update interval and combining the terms, the resulting value was basically white in nature. By expanding the double-

difference range term further and incorporating both the satellite and user position, Equation (3.18) became:

$$\begin{aligned}\Delta\nabla\rho_{xy}^{ab} &= \sqrt{(x_1 - x_{sv}^a)^2 + (x_2 - y_{sv}^a)^2 + (x_3 - z_{sv}^a)^2} \\ &\quad - \sqrt{(x_1 - x_{sv}^b)^2 + (x_2 - y_{sv}^b)^2 + (x_3 - z_{sv}^b)^2} \\ &\quad + (r_x^b - r_x^a) + \nu_{\Delta\nabla\rho}\end{aligned}\tag{3.19}$$

where

$$\begin{aligned}x_{sv}^a, y_{sv}^a, z_{sv}^a &= \text{ECEF coordinates for satellite } a \\ x_{sv}^b, y_{sv}^b, z_{sv}^b &= \text{ECEF coordinates for satellite } b \\ r_x^b &= \text{True range from satellite } b \text{ to receiver } x \\ r_x^a &= \text{True range from satellite } a \text{ to receiver } x\end{aligned}$$

In order to use Equation (3.19) in an EKF, it must be linearized to compute the filter gain and filter-computed error covariance. The linearization occurred by using a Taylor series expansion and first order approximation found in [33]. Thus, the partial derivative of the \mathbf{H} matrix was defined as:

$$\mathbf{H} [t_i; \hat{\mathbf{x}}(t_i^-)] \equiv \left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{d\mathbf{x}} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)}\tag{3.20}$$

The same treatment was also given to the phase measurement equation discussed in Section 2.2.4. After expanding the double-difference range term and combining the error sources the phase equation was:

$$\Delta\nabla\phi_{xy}^{ab} = \frac{1}{\lambda} [(r_y^a - r_x^a) - (r_y^b - r_x^b)] + \Delta\nabla N_{xy}^{ab} + \nu_{\Delta\nabla\phi}\tag{3.21}$$

where

$$\begin{aligned}\frac{1}{\lambda} [(r_y^a - r_x^a) - (r_y^b - r_x^b)] &= \frac{1}{\lambda} (\Delta\nabla r_{xy}^{ab}) \\ \Delta\nabla N_{xy}^{ab} &= \text{Double-difference integer ambiguity} \\ \nu_{\Delta\nabla\phi} &= \frac{1}{\lambda} (\Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab})\end{aligned}$$

Using all of the same assumptions about the noise vector, ν , from Equation (3.18), an expanded version of the Equation (3.21) is shown below with satellite and receiver ECEF positions incorporated:

$$\begin{aligned}\Delta\nabla\phi_{xy}^{ab} &= \frac{1}{\lambda} \left[\sqrt{(x_1 - x_{sv}^a)^2 + (x_2 - y_{sv}^a)^2 + (x_3 - z_{sv}^a)^2} \right. \\ &\quad \left. - \sqrt{(x_1 - x_{sv}^b)^2 + (x_2 - y_{sv}^b)^2 + (x_3 - z_{sv}^b)^2} \right] \\ &\quad + \frac{1}{\lambda} (r_x^b - r_x^a) + x_{1n} + \nu_{\Delta\nabla\phi}\end{aligned}\tag{3.22}$$

The next step in development of the \mathbf{H} matrix was taking the individual partial derivatives of the \mathbf{h} vector with respect to each of the \mathbf{x} vector states in Equations (3.19) and (3.22). The individual partial derivatives of the double-differenced code equation were:

$$\begin{aligned} \left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{dx_1} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)} &= \frac{x_1 - x_{sv}^a}{\sqrt{(x_1 - x_{sv}^a)^2 + (x_2 - y_{sv}^a)^2 + (x_3 - z_{sv}^a)^2}} \\ &- \frac{x_1 - x_{sv}^b}{\sqrt{(x_1 - x_{sv}^b)^2 + (x_2 - y_{sv}^b)^2 + (x_3 - z_{sv}^b)^2}} \\ &\equiv (e_1^a - e_1^b) \end{aligned} \quad (3.23)$$

$$\begin{aligned} \left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{dx_2} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)} &= \frac{x_2 - y_{sv}^a}{\sqrt{(x_1 - x_{sv}^a)^2 + (x_2 - y_{sv}^a)^2 + (x_3 - z_{sv}^a)^2}} \\ &- \frac{x_2 - y_{sv}^b}{\sqrt{(x_1 - x_{sv}^b)^2 + (x_2 - y_{sv}^b)^2 + (x_3 - z_{sv}^b)^2}} \\ &\equiv (e_2^a - e_2^b) \end{aligned} \quad (3.24)$$

$$\begin{aligned} \left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{dx_3} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)} &= \frac{x_3 - z_{sv}^a}{\sqrt{(x_1 - x_{sv}^a)^2 + (x_2 - y_{sv}^a)^2 + (x_3 - z_{sv}^a)^2}} \\ &- \frac{x_3 - z_{sv}^b}{\sqrt{(x_1 - x_{sv}^b)^2 + (x_2 - y_{sv}^b)^2 + (x_3 - z_{sv}^b)^2}} \\ &\equiv (e_3^a - e_3^b) \end{aligned} \quad (3.25)$$

Equations (3.23), (3.24) and (3.25) combine to form the difference between the unit line-of-sight vector from satellite a to the mobile receiver, and the unit line-of-sight vector from satellite b to the mobile receiver. For the purposes of this thesis, the unit line-of-sight vector is considered a row vector. This vector denoted, $(\mathbf{e}^a - \mathbf{e}^b)$, was comprised of the following three states:

$$(\mathbf{e}^a - \mathbf{e}^b) = \begin{bmatrix} (e_1^a - e_1^b) & (e_2^a - e_2^b) & (e_3^a - e_3^b) \end{bmatrix} \quad (3.26)$$

Using the $(\mathbf{e}^a - \mathbf{e}^b)$ notation, a single row in the \mathbf{H} matrix for the double-difference code measurement between satellites a and b was:

$$\mathbf{H}^{ab} = \begin{bmatrix} (\mathbf{e}^a - \mathbf{e}^b) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \quad (3.27)$$

The \mathbf{H}^{ab} vector was $9 + (n - 1)$ states long with the velocity and acceleration states included, or $3 + (n - 1)$ when the filter was in position-only mode. The first three states were always the difference between the unit line-of-sight vectors for the base satellite, a , and the current satellite, b . The last $(n - 1)$ states corresponded to the integer ambiguity states.

The double-difference phase measurement partial derivatives were identical to the code with the addition of the wavelength term, $1/\lambda$, and the ambiguity state. The value of λ would depend upon the phase observables used. It could be either the L1, L2, or widelane wavelength. The phase

partial derivatives were:

$$\begin{aligned}
\left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{dx_1} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)} &= \frac{1}{\lambda} \left[\frac{x_1 - x_{sv}^a}{\sqrt{(x_1 - x_{sv}^a)^2 + (x_2 - y_{sv}^a)^2 + (x_3 - z_{sv}^a)^2}} \right] \\
&- \frac{1}{\lambda} \left[\frac{x_1 - x_{sv}^b}{\sqrt{(x_1 - x_{sv}^b)^2 + (x_2 - y_{sv}^b)^2 + (x_3 - z_{sv}^b)^2}} \right] \\
&\equiv \frac{1}{\lambda} (e_1^a - e_1^b)
\end{aligned} \tag{3.28}$$

$$\begin{aligned}
\left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{dx_2} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)} &= \frac{1}{\lambda} \left[\frac{x_2 - y_{sv}^a}{\sqrt{(x_1 - x_{sv}^a)^2 + (x_2 - y_{sv}^a)^2 + (x_3 - z_{sv}^a)^2}} \right] \\
&- \frac{1}{\lambda} \left[\frac{x_2 - y_{sv}^b}{\sqrt{(x_1 - x_{sv}^b)^2 + (x_2 - y_{sv}^b)^2 + (x_3 - z_{sv}^b)^2}} \right] \\
&\equiv \frac{1}{\lambda} (e_2^a - e_2^b)
\end{aligned} \tag{3.29}$$

$$\begin{aligned}
\left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{dx_3} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)} &= \frac{1}{\lambda} \left[\frac{x_3 - z_{sv}^a}{\sqrt{(x_1 - x_{sv}^a)^2 + (x_2 - y_{sv}^a)^2 + (x_3 - z_{sv}^a)^2}} \right] \\
&- \frac{1}{\lambda} \left[\frac{x_3 - z_{sv}^b}{\sqrt{(x_1 - x_{sv}^b)^2 + (x_2 - y_{sv}^b)^2 + (x_3 - z_{sv}^b)^2}} \right] \\
&\equiv \frac{1}{\lambda} (e_3^a - e_3^b)
\end{aligned} \tag{3.30}$$

$$\left. \frac{d\mathbf{h}[\mathbf{x}, t_i]}{dx_{1n}} \right|_{\mathbf{x}=\hat{\mathbf{x}}(t_i^-)} = 1 \tag{3.31}$$

Using the $(\mathbf{e}^a - \mathbf{e}^b)$ notation from above, a single row in the \mathbf{H} matrix for the double-difference phase measurement was:

$$\mathbf{H}^{ab} = \left[\frac{1}{\lambda} (\mathbf{e}^a - \mathbf{e}^b) \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad \dots \quad 1 \quad \dots \quad 0 \right] \tag{3.32}$$

The 1 towards the end of the vector in Equation (3.32) corresponds to the appropriate ambiguity state in the \mathbf{x} vector. The first $(n - 1)$ rows of the \mathbf{H} matrix were associated with the double-differenced code measurements, with the last $(n - 1)$ rows of the \mathbf{H} matrix associated with the double-differenced phase measurements. The complete \mathbf{H} matrix was specified by, where the

vector \mathbf{e}^1 represented the reference satellite:

$$\mathbf{H} = \begin{bmatrix} (\mathbf{e}^1 - \mathbf{e}^2) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ (\mathbf{e}^1 - \mathbf{e}^3) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ (\mathbf{e}^1 - \mathbf{e}^n) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 0 \\ \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^2) & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & \cdots & 0 \\ \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^3) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^n) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 1 \end{bmatrix} \quad (3.33)$$

In order to put the error vector, $\nu(t_i)$, into a matrix form usable by the EKF, Equation (3.15) was used to construct the measurement error covariance matrix, \mathbf{R} . The \mathbf{R} matrix basically corresponds to the accuracy of the measurements. There are five cases which needed to be identified:

- Diagonal terms

- Case 1: Variance of code measurement errors
- Case 2: Variance of phase measurement errors

- Off-diagonal (covariance) terms

- Case 3: Covariance of two different code measurement errors
- Case 4: Covariance of two different phase measurement errors
- Case 5: Cross-covariance of code and phase measurement errors

Case 1: (Code Variance) Using Equation (3.18) gave:

$$\begin{aligned} r_{\rho^{ij}, \rho^{ij}} &= E \left\{ \nu_{\Delta\nabla\rho^{ij}} \nu_{\Delta\nabla\rho^{ij}} \right\} \\ &= E \left\{ \left(\Delta\nabla T^{ij} + \Delta\nabla I^{ij} + \Delta\nabla m^{ij} + \Delta\nabla v^{ij} \right)^2 \right\} \end{aligned} \quad (3.34)$$

Since the baseline distance between the mobile and reference receivers in this research was short (less than 5 km, typically less than 100 m), the double-differenced tropospheric, $\Delta\nabla T^{ij}$, and ionospheric, $\Delta\nabla I^{ij}$, errors were negligible and were ignored. All that remained were the multipath and noise errors. Based on the airborne testing environment, the multipath and noise errors were assumed to be completely uncorrelated and zero-mean. With these assumptions, the code variance

equation became:

$$\begin{aligned}
r_{\rho^{ij}, \rho^{ij}} &= E \left\{ (\Delta \nabla m^{ij} + \Delta \nabla v^{ij})^2 \right\} \\
&= E \left\{ \left[m_{mob}^i - m_{ref}^i - (m_{mob}^j - m_{ref}^j) + v_{mob}^i - v_{ref}^i - (v_{mob}^j - v_{ref}^j) \right]^2 \right\} \\
&= 4 E \{ m^2 + v^2 \}
\end{aligned} \tag{3.35}$$

Case 2: (Phase Variance) Case 2 was identical to case 1, the only differences were the units and magnitude of the error variance.

Case 3: (Code Covariance) Now Equation (3.34) became:

$$r_{\rho^{ij}, \rho^{xy}} = E \{ \nu_{\Delta \nabla \rho^{ij}} \nu_{\Delta \nabla \rho^{xy}} \} \tag{3.36}$$

There are two subcases which must be considered under case 3.

- Case 3A: $i \neq j \neq x \neq y$. Since the satellites were generally in different parts of the sky, the measurements from each satellite experienced different tropospheric, ionospheric, multipath and noise effects. Therefore, all code errors were assumed uncorrelated between satellites. Thus, $r_{\rho^{ij}, \rho^{xy}} = 0$, when $i \neq j \neq x \neq y$.
- Case 3B: $i = x \neq j \neq y$. The same assumptions about tropospheric and ionospheric errors from case 1 were used which reduced the problem to:

$$\begin{aligned}
r_{\rho^{ij}, \rho^{iy}} &= E \{ \nu_{\Delta \nabla \rho^{ij}} \nu_{\Delta \nabla \rho^{iy}} \} \\
&= E \{ (\Delta \nabla m^{ij} + \Delta \nabla v^{ij}) (\Delta \nabla m^{iy} + \Delta \nabla v^{iy}) \} \\
&= E \left\{ \left[m_{mob}^i - m_{ref}^i - (m_{mob}^j - m_{ref}^j) + v_{mob}^i - v_{ref}^i - (v_{mob}^j - v_{ref}^j) \right] \right. \\
&\quad \left. \left[m_{mob}^i - m_{ref}^i - (m_{mob}^y - m_{ref}^y) + v_{mob}^i - v_{ref}^i - (v_{mob}^y - v_{ref}^y) \right] \right\}
\end{aligned} \tag{3.37}$$

As in case 1, all code multipath and noise errors were completely uncorrelated. Thus, all of the j and y terms in the final part of Equation (3.37) canceled. After expanding and combining like terms:

$$r_{\rho^{ij}, \rho^{iy}} = 2 E \{ m^2 + v^2 \} = \frac{r_{\rho^{ij}, \rho^{ij}}}{2} \tag{3.38}$$

Case 4: (Phase Covariance) Case 4 was identical to case 3, the only differences were the units and magnitude of the error variance.

Case 5: (Code/Phase Cross-covariance) The base equation now became:

$$r_{\rho^{ij}, \phi^{xy}} = E \{ \nu_{\Delta \nabla \rho^{ij}} \nu_{\Delta \nabla \phi^{xy}} \} \tag{3.39}$$

Assuming that the code and phase errors were uncorrelated, every value was zero. Thus, $r_{\rho^{ij}, \phi^{xy}} = 0$. If for some reason the tropospheric errors were large, due to a long baseline distance for instance, some of the terms would be non-zero.

Finally, The complete \mathbf{R} matrix was specified by:

$$\mathbf{R} = \begin{bmatrix} r_{\rho^{ij}, \rho^{ij}} & r_{\rho^{ij}, \rho^{ik}} & \cdots & r_{\rho^{ij}, \rho^{ik}} & 0 & 0 & \cdots & 0 \\ r_{\rho^{ij}, \rho^{ik}} & r_{\rho^{ij}, \rho^{ij}} & \cdots & r_{\rho^{ij}, \rho^{ik}} & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ r_{\rho^{ij}, \rho^{ik}} & r_{\rho^{ij}, \rho^{ik}} & \cdots & r_{\rho^{ij}, \rho^{ij}} & 0 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 & r_{\phi^{ij}, \phi^{ij}} & r_{\phi^{ij}, \phi^{ik}} & \cdots & r_{\phi^{ij}, \phi^{ik}} \\ 0 & 0 & \cdots & 0 & r_{\phi^{ij}, \phi^{ik}} & r_{\phi^{ij}, \phi^{ij}} & \cdots & r_{\phi^{ij}, \phi^{ik}} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & r_{\phi^{ij}, \phi^{ik}} & r_{\phi^{ij}, \phi^{ik}} & \cdots & r_{\phi^{ij}, \phi^{ij}} \end{bmatrix} \quad (3.40)$$

where

- $r_{\rho^{ij}, \rho^{ij}}$ = double-differenced code measurement error variance
- $r_{\rho^{ij}, \rho^{ik}}$ = double-differenced code measurement error covariance
- $r_{\phi^{ij}, \phi^{ij}}$ = double-differenced phase measurement error variance
- $r_{\phi^{ij}, \phi^{ik}}$ = double-differenced phase measurement error covariance

The values used during both air and ground testing are listed in Table 3.5. The phase error numbers were based on data from Ashtech Z Surveyor receivers as given in [48]. Although not identical to the Javad® JNS100 receivers used in this research, the phase values were assumed close enough to model the phase measurement errors accurately. The code error was determined empirically through ground testing. The 12-meter code error de-weighted the code measurements, which forced the floating filter to rely more heavily on the accurate phase measurements for position estimation. Also, it helped the LAMBDA algorithm consistently return the true ambiguity set earlier in the data run.

3.4.3 Discrete-Time Models. The previous dynamics equations described a continuous-time system. Since the Kalman filter ran at periodic intervals, these equations had to be converted into discrete-time models in the following form [32]:

$$\mathbf{x}(t_{k+1}) = \Phi(t_{k+1}, t_k) \mathbf{x}(t_k) + \mathbf{w}_d \quad (3.41)$$

Table 3.5: Floating Point Filter Covariance Error Values

Term	Definition	Value
$r_{\rho^{ij}, \rho^{ij}}$	Code error variance	$4 (12 \text{ m})^2 = 576 \text{ m}^2$
$r_{\rho^{ij}, \rho^{ik}}$	Code error covariance	$2 (12 \text{ m})^2 = 288 \text{ m}^2$
$r_{\phi^{ij}, \phi^{ij}}$	Phase error variance	$4 \left(\frac{0.03 \text{ m}}{\lambda_{L1}} \right)^2 = 0.0994 \text{ cycles}^2$
$r_{\phi^{ij}, \phi^{ik}}$	Phase error covariance	$2 \left(\frac{0.03 \text{ m}}{\lambda_{L1}} \right)^2 = 0.0497 \text{ cycles}^2$

where

$$\begin{aligned}
 E \{ \mathbf{w}_d \} &= \mathbf{0} \\
 E \{ \mathbf{w}_d(t_k) \mathbf{w}_d^T(t_k) \} &= \mathbf{Q}_d \\
 E \{ \mathbf{w}_d(t_j) \mathbf{w}_d^T(t_k) \} &= \mathbf{0}, \quad t_j \neq t_k
 \end{aligned} \tag{3.42}$$

The state transition matrix, $\Phi(\Delta t)$, was calculated by:

$$\Phi(t_{k+1}, t_k) = \Phi(\Delta t) = e^{\mathbf{F} \Delta t} \tag{3.43}$$

where $\Delta t \equiv t_{k+1} - t_k$, resulting in the matrix:

$$\Phi(\Delta t) = \begin{bmatrix} 1 & 0 & 0 & \Delta t & 0 & 0 & A & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & 0 & \Delta t & 0 & 0 & A & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & 0 & \Delta t & 0 & 0 & A & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & B & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & B & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & B & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & C & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & C & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & C & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & 1 \end{bmatrix} \tag{3.44}$$

where

$$\begin{aligned}
A &= T_a^2 (e^{-\Delta t/T_a} - 1) + T_a \Delta t \\
B &= T_a (1 - e^{-\Delta t/T_a}) \\
C &= e^{-\Delta t/T_a}
\end{aligned}$$

The above $\Phi(\Delta t)$ matrix was a square matrix of $9 + (n - 1)$ states and was only used when the velocity and acceleration states were selected via the initialization file. In position-only mode, the $\Phi(\Delta t)$ matrix was a $3 + (n - 1)$ identity matrix since the \mathbf{F} matrix was all zeros.

In the general case, the discrete dynamics driving noise matrix was calculated by [32]:

$$\mathbf{Q}_d(t_k) = \int_{t_k}^{t_{k+1}} \Phi(t_{k+1}, \tau) \mathbf{G}(\tau) \mathbf{Q}(\tau) \mathbf{G}^T(\tau) \Phi^T(t_{k+1}, \tau) d\tau \quad (3.45)$$

Since the Kalman gain matrix, \mathbf{G} , was an identity matrix, Equation (3.45) can be reduced to:

$$\mathbf{Q}_d = \int_0^{\Delta t} \Phi(\tau) \mathbf{Q} \Phi^T(\tau) d\tau \quad (3.46)$$

resulting in the matrix:

$$\mathbf{Q}_d = \begin{bmatrix} D & 0 & 0 & E & 0 & 0 & G & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & D & 0 & 0 & E & 0 & 0 & G & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & D & 0 & 0 & E & 0 & 0 & G & 0 & 0 & \cdots & 0 \\ E & 0 & 0 & K & 0 & 0 & L & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & E & 0 & 0 & K & 0 & 0 & L & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & E & 0 & 0 & K & 0 & 0 & L & 0 & 0 & \cdots & 0 \\ G & 0 & 0 & L & 0 & 0 & M & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & G & 0 & 0 & L & 0 & 0 & M & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & G & 0 & 0 & L & 0 & 0 & M & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & U & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & U & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & U \end{bmatrix} \quad (3.47)$$

where

$$\begin{aligned}
D &= \frac{1}{2} T_a^5 q_a (1 - e^{-2\Delta t/T_a}) + T_a^4 q_a \Delta t (1 - 2e^{-\Delta t/T_a}) - T_a^3 q_a (\Delta t)^2 + \frac{1}{3} T_a^2 q_a (\Delta t)^3 \\
E &= T_a^4 q_a \left(\frac{1}{2} e^{-2\Delta t/T_a} - e^{-\Delta t/T_a} + \frac{1}{2} \right) + T_a^3 q_a \Delta t (e^{-\Delta t/T_a} - 1) + \frac{1}{2} T_a^2 q_a (\Delta t)^2 \\
G &= \frac{1}{2} T_a^3 q_a (1 - e^{-2\Delta t/T_a}) - T_a^2 q_a \Delta t e^{-\Delta t/T_a} \\
K &= \frac{1}{2} T_a^3 q_a \left(-e^{-2\Delta t/T_a} + 4e^{-\Delta t/T_a} + 2\frac{\Delta t}{T_a} - 3 \right) \\
L &= -\frac{1}{2} T_a^2 q_a (-e^{-2\Delta t/T_a} + 2e^{-\Delta t/T_a} - 1) \\
M &= -\frac{1}{2} T_a q_a (e^{-2\Delta t/T_a} - 1) \\
U &= q_N \Delta t
\end{aligned}$$

Just as with the Φ matrix, \mathbf{Q}_d was a $9 + (n - 1)$ square matrix when the filter had the relative velocity and acceleration states. If the filter was in position-only mode, then the \mathbf{Q}_d matrix was $3 + (n - 1)$ states defined by:

$$\mathbf{Q}_d = \begin{bmatrix} T & 0 & 0 & 0 & 0 & \cdots & 0 \\ 0 & T & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & T & 0 & 0 & \cdots & 0 \\ 0 & 0 & 0 & U & 0 & \cdots & 0 \\ 0 & 0 & 0 & 0 & U & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & U \end{bmatrix} \quad (3.48)$$

where

$$\begin{aligned}
T &= q_p \Delta t \\
U &= q_N \Delta t
\end{aligned}$$

(Note: T_a, q_a, q_p and q_N from Equations (3.47) and (3.48) are described in Table 3.3.)

That concludes the discussion on the design of the floating point Kalman filter. Now that the Kalman filter is complete, estimates of the double difference ambiguities and their associated covariance matrix can be used to generate candidate ambiguity sets.

3.5 Candidate Ambiguity Set Generation

The candidate ambiguity sets were generated using a three-step process. First, the floating point ambiguity estimates were changed using the Z-transform. Next, a discrete

search was performed over the ellipsoidal region described as the ambiguity search ellipsoid. Finally, an inverse Z-transform was performed to bring the integer ambiguity estimates back into the original measurement space. The LAMBDA routine used in this thesis was taken directly from the Delft University of Technology website in the Geodetic section. It was converted from MATLAB[®] into C++ by the author. A detailed description of the LAMBDA process and associated MATLAB[®] code can be found in [23–25].

The LAMBDA code had the option to select the top m candidate sets from the search process. This parameter was set using the initialization file. For the purposes of this research, the top two candidate ambiguity sets were selected. Based on this decision, the floating point filter was given a specified amount of time to converge to the correct floating ambiguity values. As previously stated, a wait time of two minutes was selected for the filter convergence time. Two minutes gave the best trade-off between minimal amount of delay and necessary convergence time to a floating point solution, which when sent to the LAMBDA function, generally yielded the correct set in the top two candidate sets. Of course, the number of satellites in view played a very significant role in the LAMBDA function’s ability to determine the correct integer ambiguity set. The two-minute convergence time assumed L1 phase observables. If the filter had widelane phase observables available, no convergence time was needed.

The algorithm could have a total of 50 available candidate ambiguity sets at one time. However, all 50 were not necessarily present. Each time the LAMBDA function was run, the top two ambiguity sets returned were checked against those already in the candidate ambiguity bank. If there was a new set, it was added to the bank. If there were already 50 candidate sets in the bank, the new set replaced the set with the highest minimum indicator value (i.e., the “worst” set). The minimum indicator is described in Section 3.7. During testing, if there were at least seven to eight satellites in view when the floating filter initialized, usually a maximum of about fifteen unique candidate ambiguity sets were found during an hour-long data run. Of the fifteen candidates sets, the correct set was usually in one of the top two positions.

The correct integer ambiguity set was the key to unlocking the power of the high accuracy phase measurements. There was no guarantee that the correct set would be found

in the top two sets returned from the LAMBDA function. The floating point filter was designed to help enhance the chances that the LAMBDA function would successfully return the correct set. A good discussion of the candidate ambiguity reliability can be found in [27]. However, once the correct set was determined, the desired centimeter-level accuracy was attainable (assuming no cycle slips occurred).

3.6 Multiple Model Adaptive Estimator Design

The MMAE portion of this thesis closely follows Henderson’s previous development [20]. Each elemental filter was also very similar to the floating point filter described previously. The theory for the MMAE was laid out in Section 2.5, and that theory was put into practice during the MMAE design. Lastly, filter pruning logic was added to help increase the accuracy of the MMAE blended solution and reduce the computational load.

3.6.1 Elemental Filter Design. Each elemental filter (EF) followed the form of the floating filter minus the last $(n - 1)$ ambiguity states. If the floating filter was in position-only mode, then the elemental filter’s state vector was only three states, namely the X, Y and Z ECEF relative positions. If the floating filter was in the velocity/acceleration mode, then each elemental filter had nine states. The elemental filters also followed the same FOGMA model from before. Each elemental filter was identical with respect to the system dynamics, process noise and tuning values. The only difference between the elemental filters was the incoming observables and their associated covariance matrix, since the filters started at different times. Each filter assumed its current candidate ambiguity set was the correct set and automatically adjusted the double-differenced phase observables accordingly. The corrected phase measurements were used to generate new state values and their associated residuals. The residuals were then used in the conditional probability calculation that determined which candidate set was the true fixed solution.

3.6.1.1 Elemental Filter Model Equations. The elemental filter state vector, dynamics matrix and dynamics driving noise matrix were identical to the floating filter except that the ambiguity states were removed.

The elemental filters were initialized by the current output of the floating point filter:

$$\begin{aligned}\mathbf{x}_k(t_0) &= \mathbf{x}_{(1-9)}(t_n) \\ \mathbf{P}_k(t_0) &= \mathbf{P}_{(1-9) \times (1-9)}(t_n)\end{aligned}\tag{3.49}$$

where t_n was the time that the elemental filter started with respect to the floating filter start time.

3.6.1.2 Elemental Filter Measurement Model. The elemental filters used an almost identical measurement model to the floating point filter. The biggest difference was the removal of all code measurements. Due to the noisy nature of the code measurements, they were precluded from the MMAE filters. The code measurements helped the floating filter with a “ball park” starting point, and prevented phase errors from ruining the floating point solution. Also, the phase measurements were pre-corrected for a particular ambiguity set prior to incorporation. The measurement vector was $(n - 1)$ states given by:

$$\mathbf{z}(t_i) = \begin{bmatrix} (\Delta \nabla \phi_{xy}^{12})' & (\Delta \nabla \phi_{xy}^{13})' & \cdots & (\Delta \nabla \phi_{xy}^{1n})' \end{bmatrix}\tag{3.50}$$

where

$$\begin{aligned}(\Delta \nabla \phi_{xy}^{12})' &= \text{ambiguity-corrected double-differenced phase measurement} \\ &\quad \text{between satellites 1 and 2} \\ (\Delta \nabla \phi_{xy}^{13})' &= \text{ambiguity-corrected double-differenced phase measurement} \\ &\quad \text{between satellites 1 and 3} \\ &\vdots \\ (\Delta \nabla \phi_{xy}^{1n})' &= \text{ambiguity-corrected double-differenced phase measurement} \\ &\quad \text{between satellites 1 and } n\end{aligned}$$

The computation of the ambiguity corrected phase measurements was defined by:

$$(\Delta \nabla \phi_{xy}^{ab})' = \Delta \nabla \phi_{xy}^{ab} - \Delta \nabla \tilde{N}_{xy}^{ab}\tag{3.51}$$

where

$$\begin{aligned}
\Delta\nabla\phi_{xy}^{ab} &= \text{the raw double-differenced phase measurement} \\
&\quad \text{between satellites } a \text{ and } b \\
\Delta\nabla\tilde{N}_{xy}^{ab} &= \text{the double-differenced candidate ambiguity} \\
&\quad \text{value between satellites } a \text{ and } b
\end{aligned}$$

A candidate ambiguity set was made up of $(n - 1)$ candidate ambiguity terms from Equation (3.51). Assuming no cycle slips or loss-of-locks, the $\Delta\nabla\tilde{N}_{xy}^{ab}$ term for a particular satellite set would remain constant. Each elemental filter had its own unique candidate ambiguity set. There may have been repeats in individual terms, but at least one $\Delta\nabla\tilde{N}_{xy}^{ab}$ term was different for each set. This was the main difference between the elemental filters. Also, since the elemental filters may have started at different times during a data run, their associated covariance matrices could vary.

Combining Equations (3.21) and (3.51) gave:

$$\left(\Delta\nabla\phi_{xy}^{ab}\right)' = \frac{1}{\lambda} \left(\Delta\nabla r_{xy}^{ab} + \Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab} \right) + \Delta\nabla N_{xy}^{ab} - \Delta\nabla\tilde{N}_{xy}^{ab} \quad (3.52)$$

As discussed before in Section 3.4.2, the measurement model must be in the form:

$$\mathbf{z}(t_i) = \mathbf{h}[\mathbf{x}(t_i), t_i] + \mathbf{v}(t_i) \quad (3.53)$$

Expanding the range term, and combining the others from Equation (3.52) gave:

$$\left(\Delta\nabla\phi_{xy}^{ab}\right)' = \frac{1}{\lambda} \left[r_y^a - r_x^a - (r_y^b - r_x^b) \right] + \nu_{\Delta\nabla\phi} \quad (3.54)$$

where

$$\begin{aligned}
\Delta\nabla r_{xy}^{ab} &= \frac{1}{\lambda} \left[r_y^a - r_x^a - (r_y^b - r_x^b) \right] \\
\nu_{\Delta\nabla\phi} &= \frac{1}{\lambda} \left(\Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab} \right) + \Delta\nabla N_{xy}^{ab} - \Delta\nabla\tilde{N}_{xy}^{ab}
\end{aligned}$$

As with the floating filter, the $\nu_{\Delta\nabla\phi}$ vector was modeled as white noise and the same assumptions apply. There was one big difference however. The $\Delta\nabla N_{xy}^{ab} - \Delta\nabla\tilde{N}_{xy}^{ab}$ term was not necessarily “white”. In fact, the only time this term was “white”, was when the candidate ambiguity set was the true ambiguity set. When this happened, $\Delta\nabla N_{xy}^{ab} - \Delta\nabla\tilde{N}_{xy}^{ab} = 0$. For any other ambiguity set there would always be a constant bias in the term. This fact

was leveraged in the residual calculations. The constant bias would increase the residuals, making the true set stand out in comparison. There was an exception to this rule: when a particular, but wrong, candidate ambiguity set actually “fit” the measurements better than the true set. In this case the wrong set’s residuals would be lower than the true ambiguity set’s residuals. When this occurred, the blended solution from the MMAE would be pulled further away from the true relative position. The stronger the fit of the incorrect ambiguity set, the further away the blended solution would be from the true solution [20]. This property of the MMAE blended solution prompted the follow-on research in this thesis.

Continuing the MMAE development, the measurement model must be linearized for use in the filter. The linearization process was the same as the floating filter discussed in Section 3.4.3. However, the final \mathbf{H} matrix is shown below since it did vary slightly from the floating filter.

When the floating filter was in velocity/acceleration mode, a single row in MMAE \mathbf{H} matrix was:

$$\mathbf{H}^{ab} = \begin{bmatrix} \frac{1}{\lambda} (\mathbf{e}_y^a - \mathbf{e}_y^b) & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (3.55)$$

If the floating filter was in position-only mode, then a single row became:

$$\mathbf{H}^{ab} = \begin{bmatrix} \frac{1}{\lambda} (\mathbf{e}_y^a - \mathbf{e}_y^b) \end{bmatrix} \quad (3.56)$$

where

$$\begin{aligned} \mathbf{e}_y^a &= \begin{bmatrix} e_1^a & e_2^a & e_3^a \end{bmatrix} = \text{unit line-of-sight vector between receiver and satellite } a. \\ \mathbf{e}_y^b &= \begin{bmatrix} e_1^b & e_2^b & e_3^b \end{bmatrix} = \text{unit line-of-sight vector between receiver and satellite } b. \end{aligned}$$

The full MMAE \mathbf{H} matrix was the last $(n - 1)$ rows of the floating filter \mathbf{H} matrix shown in Equation (3.33), without the ones for the ambiguity states. The full MMAE \mathbf{H} matrix when in velocity/acceleration mode was an $(n - 1)$ -by-9 matrix specified by:

$$\mathbf{H} = \begin{bmatrix} \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^2) & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^3) & 0 & 0 & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^n) & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (3.57)$$

In position-only mode:

$$\mathbf{H} = \begin{bmatrix} \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^2) \\ \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^3) \\ \vdots \\ \frac{1}{\lambda}(\mathbf{e}^1 - \mathbf{e}^n) \end{bmatrix} \quad (3.58)$$

The measurement error covariance matrix \mathbf{R} was identical to the last $(n - 1)$ rows and columns of the floating filter \mathbf{R} matrix. The MMAE \mathbf{R} matrix was still defined by Equation (3.15). The full MMAE \mathbf{R} matrix was specified by:

$$\mathbf{R} = \begin{bmatrix} r_{\phi^{ij}, \phi^{ij}} & r_{\phi^{ij}, \phi^{ik}} & \cdots & r_{\phi^{ij}, \phi^{ik}} \\ r_{\phi^{ij}, \phi^{ik}} & r_{\phi^{ij}, \phi^{ij}} & \cdots & r_{\phi^{ij}, \phi^{ik}} \\ \vdots & \vdots & \ddots & \vdots \\ r_{\phi^{ij}, \phi^{ik}} & r_{\phi^{ij}, \phi^{ik}} & \cdots & r_{\phi^{ij}, \phi^{ij}} \end{bmatrix} \quad (3.59)$$

where

$r_{\phi^{ij}, \phi^{ij}}$ = double-differenced phase measurement error variance

$r_{\phi^{ij}, \phi^{ik}}$ = double-differenced phase measurement error covariance

The values for $r_{\phi^{ij}, \phi^{ij}}$ and $r_{\phi^{ij}, \phi^{ik}}$ were the same as the floating filter shown in Table 3.5.

3.6.1.3 Elemental Filter Discrete-Time Models. Just as with the floating filter, the elemental filter must use discrete-time models for implementation in an EKF construct. The method described in Section 3.4.3 was used for the elemental filters, once

again assuming $\Delta t \equiv t_{k+1} - t_k$.

$$\Phi(\Delta t) = \begin{bmatrix} 1 & 0 & 0 & \Delta t & 0 & 0 & A_e & 0 & 0 \\ 0 & 1 & 0 & 0 & \Delta t & 0 & 0 & A_e & 0 \\ 0 & 0 & 1 & 0 & 0 & \Delta t & 0 & 0 & A_e \\ 0 & 0 & 0 & 1 & 0 & 0 & B_e & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & B_e & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & B_e \\ 0 & 0 & 0 & 0 & 0 & 0 & C_e & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & C_e & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & C_e \end{bmatrix} \quad (3.60)$$

where

$$\begin{aligned} A_e &= T_a^2 (e^{-\Delta t/T_a} - 1) + T_a \Delta t \\ B_e &= T_a (1 - e^{-\Delta t/T_a}) \\ C_e &= e^{-\Delta t/T_a} \end{aligned}$$

The above $\Phi(\Delta t)$ matrix was a square matrix of nine states and was only used when the velocity and acceleration states were turned on in the floating filter. In position-only mode, the $\Phi(\Delta t)$ matrix was a 3-by-3 identity matrix since the \mathbf{F} matrix was all zeros.

The discrete driving noise matrix was defined by:

$$\mathbf{Q}_d = \begin{bmatrix} D_e & 0 & 0 & E_e & 0 & 0 & G_e & 0 & 0 \\ 0 & D_e & 0 & 0 & E_e & 0 & 0 & G_e & 0 \\ 0 & 0 & D_e & 0 & 0 & E_e & 0 & 0 & G_e \\ E_e & 0 & 0 & K_e & 0 & 0 & L_e & 0 & 0 \\ 0 & E_e & 0 & 0 & K_e & 0 & 0 & L_e & 0 \\ 0 & 0 & E_e & 0 & 0 & K_e & 0 & 0 & L_e \\ G_e & 0 & 0 & L_e & 0 & 0 & M_e & 0 & 0 \\ 0 & G_e & 0 & 0 & L_e & 0 & 0 & M_e & 0 \\ 0 & 0 & G_e & 0 & 0 & L_e & 0 & 0 & M_e \end{bmatrix} \quad (3.61)$$

where

$$\begin{aligned}
D_e &= \frac{1}{2} T_a^5 q_a (1 - e^{-2\Delta t/T_a}) + T_a^4 q_a \Delta t (1 - 2e^{-\Delta t/T_a}) - T_a^3 q_a (\Delta t)^2 + \frac{1}{3} T_a^2 q_a (\Delta t)^3 \\
E_e &= T_a^4 q_a (\frac{1}{2} e^{-2\Delta t/T_a} - e^{-\Delta t/T_a} + \frac{1}{2}) + T_a^3 q_a \Delta t (e^{-\Delta t/T_a} - 1) + \frac{1}{2} T_a^2 q_a (\Delta t)^2 \\
G_e &= \frac{1}{2} T_a^3 q_a (1 - e^{-2\Delta t/T_a}) - T_a^2 q_a \Delta t e^{-\Delta t/T_a} \\
K_e &= \frac{1}{2} T_a^3 q_a (-e^{-2\Delta t/T_a} + 4e^{-\Delta t/T_a} + 2\frac{\Delta t}{T_a} - 3) \\
L_e &= -\frac{1}{2} T_a^2 q_a (-e^{-2\Delta t/T_a} + 2e^{-\Delta t/T_a} - 1) \\
M_e &= -\frac{1}{2} T_a q_a (e^{-2\Delta t/T_a} - 1)
\end{aligned}$$

Just as with the Φ matrix, the \mathbf{Q}_d was a 9-by-9 square matrix when the filter has the velocity and acceleration states. If the filter was in position-only mode, then the \mathbf{Q}_d matrix was 3-by-3 states defined by:

$$\mathbf{Q}_d = \begin{bmatrix} T_e & 0 & 0 \\ 0 & T_e & 0 \\ 0 & 0 & T_e \end{bmatrix} \quad (3.62)$$

where

$$T_e = q_p \Delta t$$

(Note: T_a , q_a and q_p from Equations (3.61) and (3.62) are described in Table 3.3)

3.6.2 Conditional Probability Calculations. The heart of the MMAE calculations was the conditional probability. The probability assigned to each candidate ambiguity set was used to determine the weighted or blended MMAE position solution. An overview of the conditional probability is shown in Figure 3.4 and described in detail in Section 2.5.1. For clarification the probability equations are repeated below:

$$p_k(t_i) = \frac{f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i|\mathbf{a}_k, \mathbf{Z}_{i-1}) p_k(t_{i-1})}{\sum_{j=1}^k f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i|\mathbf{a}_j, \mathbf{Z}_{i-1}) p_j(t_{i-1})} \quad (3.63)$$

where

$$\begin{aligned}
f_{\mathbf{z}(t_i)|\mathbf{a}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i|\mathbf{a}_k, \mathbf{Z}_{i-1}) &= \frac{1}{(2\pi)^{m/2} \sqrt{|\mathbf{A}_k(t_i)|}} e^{\{\bullet\}} \\
\{\bullet\} &= -\frac{1}{2} \mathbf{r}_k^T(t_i) \mathbf{A}_k^{-1}(t_i) \mathbf{r}_k(t_i)
\end{aligned} \quad (3.64)$$

and

$$\begin{aligned}
\mathbf{A}(t_i) &= \mathbf{H}(t_i) \mathbf{P}(t_i^-) \mathbf{H}^T(t_i) + \mathbf{R}(t_i) \\
\mathbf{r}(t_i) &= \mathbf{z}_i - \mathbf{h}[\hat{\mathbf{x}}(t_i^-), t_i]
\end{aligned} \quad (3.65)$$

Originally, Henderson [20] used the updated state estimate and covariance after measurement incorporation, called “post-fit” residuals, to perform the probability calculations. It was thought that the updated residuals would be less noisy and therefore provide a more accurate probability estimate. However, further research performed by Ormsby [44], proved that there was no theoretical difference between the “pre-fit” and “post-fit” residuals for this type of MMAE. To test the theory proposed by Ormsby, both “pre-fit” and “post-fit” residual calculations were performed during test runs. The results are shown in Chapter IV.

Also, since every elemental filter in Henderson’s thesis started at the same time, the leading $\frac{1}{(2\pi)^{m/2}\sqrt{|\mathbf{A}_k(t_i)|}}$ coefficient was removed during his testing [20]. For this thesis, each elemental filter did not necessarily start at the same time, so the coefficient must be included.

3.6.2.1 Lower Probability Bounds. The lower probability bound, ε , described in Section 2.5.1.1 was set at 0.001. The lower probability bound prevented the MMAE from assigning a value of zero to a particular conditional probability set. Although the probability of a particular elemental filter may be nonzero, the elemental filter’s conditional probability would remain at zero if it were ever computed as zero, based on the numerator in Equation (3.63). Also, to improve the weighted MMAE solution’s accuracy, only candidate ambiguity sets with probabilities above 1 percent were included in the blended solution. However, care was taken after any candidate set was removed ensuring the remaining probabilities were re-weighted based on a total density of 100 percent.

3.6.3 Elemental Filter Pruning. Pruning elemental filters from the MMAE served a dual purpose. First, it prevented an erroneous candidate set from degrading the MMAE accuracy. Second, it significantly reduced the computational load. Along with all of the required MMAE calculations on a candidate set, the minimum indicator calculations also needed to be performed. Within the minimum indicator calculations (described below in Section 3.7) was a second order least squares fit on the time history of the sum-of-squared (SOS) residuals. The least squares fit was computationally expensive, especially as the number of data points to be fit increased.

The method used to prune elemental filters was based on the floating filter’s relative position standard deviation. If the difference between a candidate ambiguity set’s relative

position and the floating filter's relative position was outside of three standard deviations (according to the floating filter estimation), the associated elemental filter was selected for removal.

It should be noted that most MMAEs were designed to look for changes within a particular system. This MMAE was no different. It also looked for changes in the phase measurements that best matched the ambiguity candidates. However, since each elemental filter represented a constant bias, namely a particular ambiguity set, once that set was determined to be incorrect, in all likelihood it would not be the correct set at any time during the remainder of a test run. But, there was a possibility that one of the incorrect sets could become the correct set based on certain errors.

There was always the possibility of cycle slips or loss-of-lock during a data run. Therefore, incorrect ambiguity sets in the candidate bank could potentially be the correct set at a later time, if a cycle slip or loss-of-lock occurred. If a candidate set was within the previously described three-sigma value, it was best to have the set in the candidate bank as a possible solution should a cycle slip occur. The algorithm only used candidate sets which had a conditional probability above 1 percent in the weighted position calculation. Thus, even though there might have been numerous incorrect candidate ambiguity sets in the bank, this usually would not grossly affect the blended relative position solution. The exception to this rule was when an incorrect ambiguity set "fit" the measurements better than the true ambiguity set, as discussed previously.

Finally, in order for the MMAE solution to be a true fixed-integer solution, a single set must absorb all of the elemental probability. This, of course, cannot happen since the algorithm artificially injected a lower probability bound on any candidate set which was assigned a probability value lower than 0.001. However, in practice, one ambiguity set would usually absorb almost all of the probability, leaving a very small amount to be divided among the remaining candidate ambiguity sets. When this happened, the 1 percent weighted position removal logic precluded every set except the "high" probability set from being used in the relative position calculation. Now the MMAE "blended" solution only contained one ambiguity set, so it was a true fixed-integer solution. Of course, the "high" probability set was not necessarily the correct ambiguity set. In this case, the navigation

solution was inaccurate, but based on the three-sigma pruning logic, the solution was still reasonable.

3.7 *Minimum Indicator Design*

The major thrust of this thesis was the design and development of the new minimum indicator variable for candidate ambiguity set verification. Certain characteristics of the candidate ambiguity sets lend themselves to identifying the true ambiguity set properly over time. First, conditional residuals must be calculated. There would be $(n - 1)$ conditional residuals, one for each individual ambiguity candidate in a particular ambiguity set. Next, the SOS residuals for each candidate set was calculated. The time history of SOS residuals for each candidate set was saved and used in a second order least squares fit to determine the coefficients on the x^2 and x terms. Finally, a combination of the average SOS residuals and average coefficients was used to calculate the minimum indicator value.

3.7.1 Conditional Residuals Calculations. The first and most important part of the minimum indicator section was the conditional residuals calculation. The conditional residuals were the residuals conditioned on the fact that a particular candidate ambiguity set was the true set. Unlike the residuals that were calculated by the MMAE elemental filters, the conditional residuals were an epoch-by-epoch calculation with no knowledge of previous measurements or residuals. The conditional residuals could be likened to a snapshot in time, and were calculated for each candidate ambiguity set.

Before the conditional residual calculation was performed, a small correction was applied to the raw double-differenced phase measurements. The correction term was due to small differences in receiver clock errors. The receiver clock was used to determine the transmit time for the GPS signal, which in turn determined satellite positions. The difference in clock errors meant each receiver calculated a slightly different satellite position for the same satellite at a particular GPS time.

The difference between the calculated satellite positions was specified by:

$$\begin{aligned}\Delta \mathbf{p}^a &= \mathbf{p}_x^a - \mathbf{p}_y^a \\ \Delta \mathbf{p}^b &= \mathbf{p}_x^b - \mathbf{p}_y^b\end{aligned}\tag{3.66}$$

where

\mathbf{p}_x^a = ECEF position for satellite a from receiver x

\mathbf{p}_y^a = ECEF position for satellite a from receiver y

\mathbf{p}_x^b = ECEF position for satellite b from receiver x

\mathbf{p}_y^b = ECEF position for satellite b from receiver y

The dot product of the delta terms and the satellite unit line-of-sight vectors form the correction term:

$$\frac{1}{\lambda} \left[\mathbf{e}^a \cdot \Delta \mathbf{p}^a - \mathbf{e}^b \cdot \Delta \mathbf{p}^b \right] \quad (3.67)$$

where

λ = wavelength (L1, L2, or widelane)

\mathbf{e}^a = unit line-of-sight vector for satellite a

\mathbf{e}^b = unit line-of-sight vector for satellite b

Finally, the double-difference phase was corrected by:

$$(\Delta \nabla \phi)' = \Delta \nabla \phi - \frac{1}{\lambda} \left[\mathbf{e}^a \cdot \Delta \mathbf{p}^a - \mathbf{e}^b \cdot \Delta \mathbf{p}^b \right] \quad (3.68)$$

where

$\Delta \nabla \phi$ = raw double-differenced phase measurements

Now that the double-differenced phase was corrected for receiver clock errors, a least squares approach was used to calculate the conditional residuals. First, the measurement vector, \mathbf{z} , was defined as:

$$\mathbf{z} = (\Delta \nabla \phi)' \quad (3.69)$$

The measurements were modified by the candidate ambiguity sets just as with the elemental filters in the MMAE by:

$$\mathbf{z}_n = \mathbf{z} - \Delta \nabla \tilde{N}_n \quad (3.70)$$

The $\Delta \nabla \tilde{N}_n$ term was a particular candidate ambiguity set. Now each candidate ambiguity set was applied to the corrected double-differenced phase measurements forming K unique measurement vectors, \mathbf{z}_n (K was the number of elemental filters, or equivalently, the number

of candidate ambiguity sets). Using Equation (3.14) as a base and the \mathbf{H} matrix development in the same section, the measurement vector was described by:

$$\mathbf{z}_n = \mathbf{H} \Delta \mathbf{x} \quad (3.71)$$

The \mathbf{H} matrix in Equation (3.71) was identical to the one described in Section 3.6.1.2. The state estimate was calculated by rearranging Equation (3.71). Since the \mathbf{H} matrix was not square, a pseudo-inverse was performed using:

$$\Delta \hat{\mathbf{x}} = \left[(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \right] \mathbf{z}_n \quad (3.72)$$

Looking back, it would have been more appropriate to include the measurement error covariance matrix, \mathbf{R} , in Equation (3.72). However, the results presented in Chapter IV should not be significantly affected by its removal. From the development in Section 2.4.5 the residual calculation was performed by:

$$\mathbf{r} = \mathbf{z}_n - \mathbf{H} \Delta \hat{\mathbf{x}} \quad (3.73)$$

Equation (3.73) was the basis from which a direct calculation of the residuals was possible without explicitly calculating a state estimate. By combining Equations (3.72) and (3.73), the resulting residual calculation became:

$$\begin{aligned} \mathbf{r} &= \mathbf{z}_n - \mathbf{H} \left[(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \right] \mathbf{z}_n \\ \mathbf{r} &= \left[\mathbf{I} - \mathbf{H} (\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \right] \mathbf{z}_n \end{aligned} \quad (3.74)$$

The bracketed portion in the final line of Equation (3.74) was calculated once per epoch since its value did not change based on the candidate ambiguities. Finally, each candidate \mathbf{z}_n vector was applied, resulting in the residuals conditioned upon a particular ambiguity set.

Equation (3.74) allowed for a quick and efficient method of determining the residuals for each candidate ambiguity set. The SOS residuals calculation was performed by:

$$\text{SOS residuals} = (r_{12})^2 + (r_{13})^2 + (r_{14})^2 + \cdots + (r_{1n})^2 \quad (3.75)$$

where

$$\begin{aligned}
 r_{12} &= \text{residual for double-differenced phase between satellites 1 and 2} \\
 r_{13} &= \text{residual for double-differenced phase between satellites 1 and 3} \\
 r_{14} &= \text{residual for double-differenced phase between satellites 1 and 4} \\
 &\vdots \\
 r_{1n} &= \text{residual for double-differenced phase between satellites 1 and n}
 \end{aligned}$$

The minimum SOS residuals was the primary indicator for the correct ambiguity set. However, there were local minima for which a particular ambiguity set might “fit” the measurements better than the true set. As time elapsed, the incorrect ambiguity set’s SOS residuals began to increase again as seen in Figure 3.6. The time elapse was directly related to satellite geometry. Satellite geometry was the main component in determining the proper ambiguity set. The reason is apparent with the unit line-of-sight vectors included in the \mathbf{H} matrix.

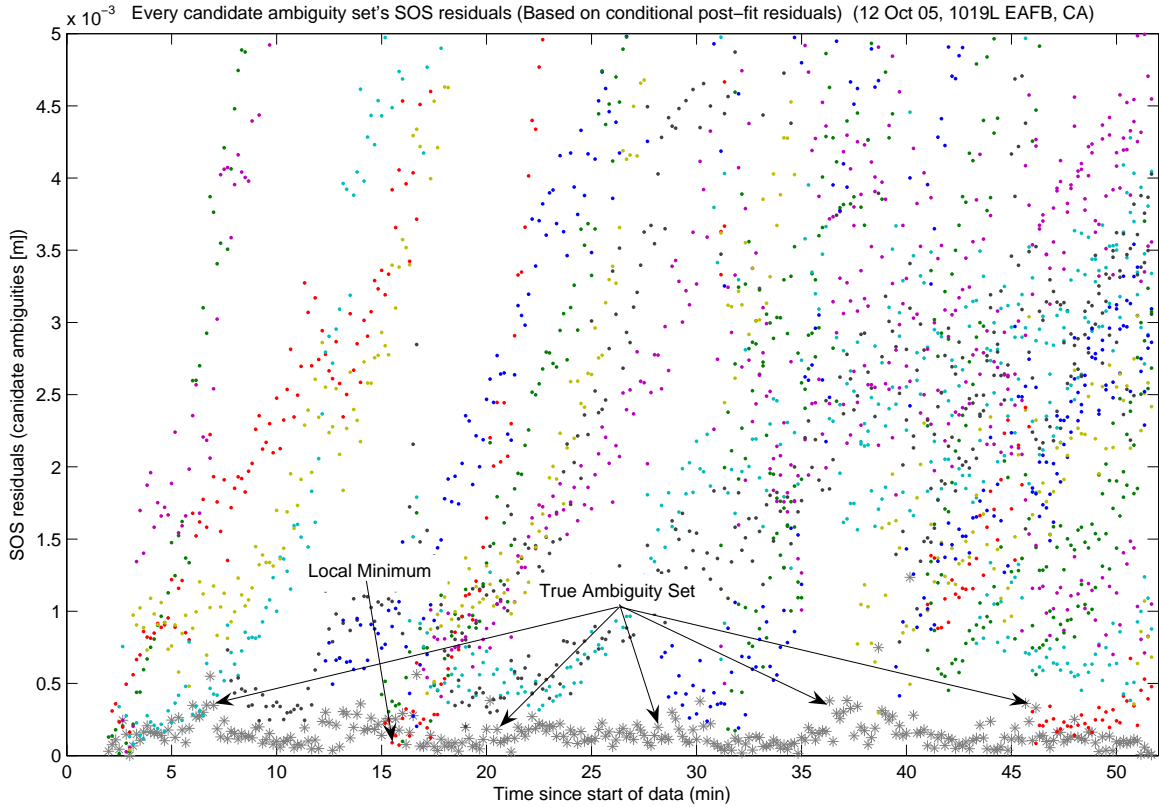


Figure 3.6: SOS Residuals Example

3.7.2 Sum-of-Squared Residual Description. As mentioned above, the primary thrust of this thesis was developing a new technique to determine the true ambiguity set based on the nature of the sum-of-squared (SOS) residuals of the least squares position estimate, after measurement incorporation. This new technique is called the minimum indicator. The minimum indicator took advantage of unique properties of the true ambiguity set's residuals.

Normally, the true ambiguity set's residuals were lower than any other candidate set. However, errors on the phase measurements, such as multipath and noise, can make an erroneous ambiguity set's measurements appear to fit the position solution better than the true set. This phenomenon normally only occurred for a short time. As satellite geometry and errors changed, the erroneous set's residuals began to increase again above those of the true ambiguity set. The time frame when the erroneous ambiguity set's residuals were below the true set was described as a local minimum. The decrease-increase cycle of an erroneous ambiguity set's residuals formed a shape similar to a parabolic curve. In contrast, the true ambiguity residuals were generally a horizontal line. Figure 3.7 shows an example of the parabolic nature of the erroneous residuals, compared with the horizontal nature of the true residuals. The data for Figure 3.7 came from the widelane data run on 28 May 2004 at the Air Force Institute of Technology (AFIT). However, the SOS residuals for L1 observables displayed the same characteristics.

The erroneous ambiguity set's residuals followed the form of $y = ax^2 + bx + c$, but the true set's residuals were of the form $y = bx + c$. Since the true residuals were horizontal, the b coefficient on the x term was theoretically zero (so were the coefficients on all other higher-order terms, including the second-order a coefficient). Also, the true ambiguity set's y-axis intercept, c , should be the smallest based on the true set's horizontal nature and low residuals.

By using a least squares fit on the residuals, the a , b , and c coefficients can be determined for each ambiguity set (assuming there are at least three data points). Based on the discussion above, the set with the smallest coefficients should, in theory, be the true ambiguity set. The minimum indicator got its namesake from the process of picking the ambiguity set with the “minimum”, or smallest, coefficients.

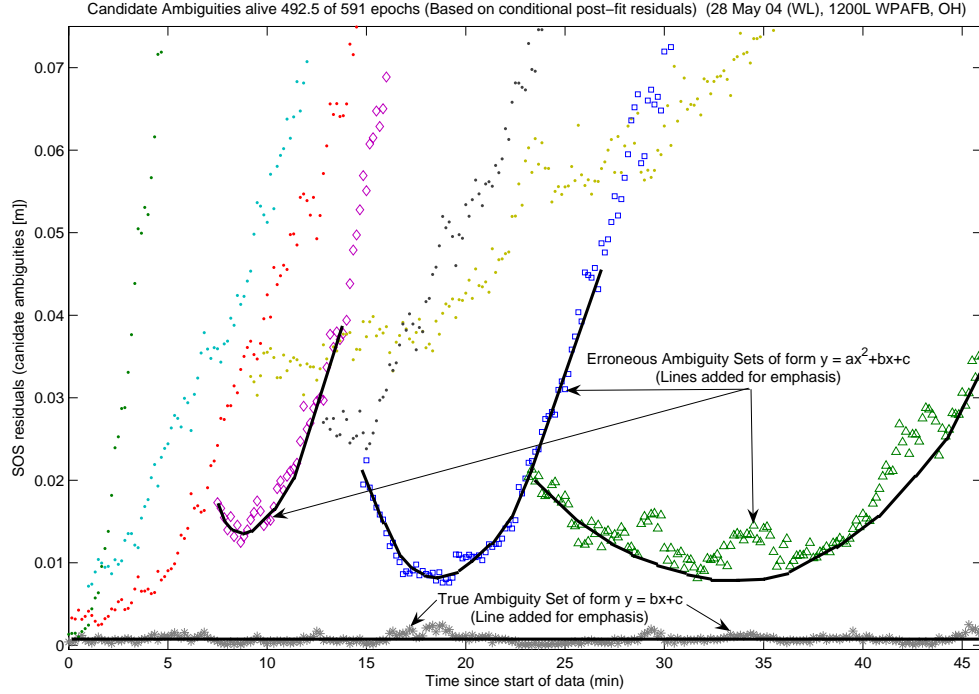


Figure 3.7: SOS Residual Comparison

3.7.3 Least Squares Fit Calculation. The previous section described in detail the special characteristics of the SOS residuals for both correct and incorrect ambiguity sets. By using the time history of SOS residuals and least squares fit equations, the algorithm provided one more piece to the minimum indicator “puzzle”. The greatest discriminator/filter ever constructed was the human eye. A human could examine Figures 3.6 and 3.7 and determine the correct ambiguity set with relative ease. The trick was to determine the elements which allow a human to identify the true ambiguity set correctly. One of those elements was the slope of the SOS residual line.

In order to reduce the number of data points collected, the SOS residuals were saved once every 10 seconds (the SOS residuals calculation was done every second). The data point reduction served a dual purpose. First, it limited the amount of required memory, and more importantly, significantly reduced the computational load of the least squares fit. During an hour-long data run, the number of data points was reduced from 3600 to only 360. Since a general idea of the shape of the residuals was desired, there was no need to collect data points at a one Hz rate.

Before the least squares fit was accomplished, a sufficient amount of data points were required. A two-minute, initial SOS residuals acquisition time was used before any least squares fit calculations were performed. Based on the SOS residuals “save interval” of ten seconds, the initial acquisition time resulted in 12 data points. This was more than enough data points to perform the calculation and determine trends in the SOS residuals.

The least squares fit needed two vectors and the specified order to accomplish the calculation. The first required vector was the actual data points, denoted as \mathbf{y} . The m -by-1 \mathbf{y} vector was specified by:

$$\mathbf{y} = \begin{bmatrix} y_1 & y_2 & y_3 & \cdots & y_m \end{bmatrix}^T \quad (3.76)$$

where

$$\begin{aligned} y_1 &= \text{the first data point} \\ y_2 &= \text{the second data point} \\ y_3 &= \text{the third data point} \\ &\vdots \\ y_m &= \text{the } m^{th} \text{ data point} \end{aligned}$$

The second vector, denoted by \mathbf{x} , was the time difference between each data point, and specified as:

$$\mathbf{x} = \begin{bmatrix} \Delta t_1 & \Delta t_2 & \Delta t_3 & \cdots & \Delta t_m \end{bmatrix}^T \quad (3.77)$$

where

$$\begin{aligned} \Delta t_1 &= t_1 - t_0 \\ \Delta t_2 &= t_2 - t_0 \\ \Delta t_3 &= t_3 - t_0 \\ &\vdots \\ \Delta t_m &= t_m - t_0 \end{aligned}$$

Since the sample rate of the SOS residuals was once every 10 seconds ($0.166\overline{6}$ minutes), the \mathbf{x} vector was an m -by-1 vector, where $\Delta t_m = 0.166\overline{6} m$ minutes.

From the \mathbf{x} vector a new matrix, \mathbf{A} , was constructed based on the desired order. It was an m -by- $(\text{order}+1)$ matrix given by:

$$\mathbf{A} = \begin{bmatrix} 1 & \Delta t_1 & (\Delta t_1)^2 & \cdots & (\Delta t_1)^{\text{order}} \\ 1 & \Delta t_2 & (\Delta t_2)^2 & \cdots & (\Delta t_2)^{\text{order}} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \Delta t_m & (\Delta t_m)^2 & \cdots & (\Delta t_m)^{\text{order}} \end{bmatrix} \quad (3.78)$$

Specifically for this thesis, the \mathbf{A} matrix was:

$$\mathbf{A} = \begin{bmatrix} 1 & 0.166\bar{6} & 0.027\bar{7} \\ 1 & 0.333\bar{3} & 0.111\bar{1} \\ \vdots & \vdots & \vdots \\ 1 & 0.166\bar{6}m & (0.166\bar{6}m)^2 \end{bmatrix} \quad (3.79)$$

The units in Equation (3.79) are minutes. The conversion from seconds to minutes was done to make the final computation easier and promote better discrimination between different ambiguity set's coefficients.

The final step was the computation of the coefficient vector, \mathbf{c} . The coefficient calculation was:

$$\mathbf{c} = \left[(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \right] \mathbf{y} \quad (3.80)$$

The \mathbf{c} vector was $(\text{order}+1)$ -by-1 with the highest coefficient value in the last location.

3.7.4 Minimum Indicator Calculation. The minimum indicator calculation used the previous development with some slight modifications to determine the correct ambiguity set. Since the SOS residuals for each ambiguity set changed from epoch to epoch, the coefficients from the least squares fit also changed slightly when the calculation was performed. Therefore, the average SOS residuals and coefficients for each particular ambiguity set was used in the minimum indicator calculation. The average SOS residuals was calculated as:

$$\overline{\text{SOS residuals}^k} = \frac{1}{m} \sum_{i=1}^m \text{SOS residuals}^k(t_i) \quad (3.81)$$

where

$$\begin{aligned} m &= \text{Number of individual SOS residuals for ambiguity set } k \\ \text{SOS residuals}^k(t_i) &= \text{SOS residuals for ambiguity set } k \text{ at epoch } i \end{aligned}$$

The minimum indicator calculation only used the c_2 and c_1 coefficients from the x^2 and x terms of the characteristic equation, $c_2 x^2 + c_1 x + c_0$. Once again, the mean values for both the c_2 and c_1 coefficients was calculated with a minor modification. Since the true ambiguity set would have a slope of zero, the c_2 and c_1 should theoretically be zero. Of course, in practice the slopes were not exactly zero, but very close. There may have been sections of the SOS residuals that had a negative slope. The real indicator of interest was the magnitude, not necessarily the direction of the slope. Therefore, the average for the coefficients was the average of their magnitude calculated by:

$$\begin{aligned} \overline{c_2^k} &= \frac{1}{l} \sum_{i=1}^l |c_2^k(t_i)| \\ \overline{c_1^k} &= \frac{1}{l} \sum_{i=1}^l |c_1^k(t_i)| \end{aligned} \tag{3.82}$$

where

$$\begin{aligned} l &= \text{Number of individual coefficients for ambiguity set } k \\ c_2^k(t_i) &= c_2 \text{ coefficient for ambiguity set } k \text{ at epoch } i \\ c_1^k(t_i) &= c_1 \text{ coefficient for ambiguity set } k \text{ at epoch } i \end{aligned}$$

As previously discussed, a wrong candidate ambiguity set may have a local minimum where its SOS residuals were lower than those of the true ambiguity set. This was the primary reason the least squares fit calculation was performed. The additional information gathered from the coefficients helped discriminate between erroneous sets and the true ambiguity set. However, there was still the possibility that a new ambiguity set returned from the LAMBDA function would be in its local minimum stage. If so, the SOS residuals and the slope would be less than or equal to the true ambiguity set's values. If the floating filter had been running for an extended time and the true ambiguity set was identified, then the algorithm should not be deceived by the new ambiguity set. To help alleviate this problem, the average SOS residuals and coefficients were divided by the total number of epochs a candidate ambiguity set was "alive". "Weighting" was increased the longer a candidate ambiguity set remained in the ambiguity bank.

The final value used in the minimum indicator calculation was the number of times a candidate set was returned from the LAMBDA function as the number one set. This value was added to the total number of epochs a candidate was alive, which increased the “weight” of a particular candidate ambiguity set. Finally, the minimum indicator calculation became:

$$\text{Minimum Indicator}^k = \frac{\overline{\text{SOS residuals}^k} + \overline{c_2^k} + \overline{c_1^k}}{m + i} \quad (3.83)$$

where

- $m =$ Total number of epochs ambiguity set k has been alive
- $i =$ Number of times ambiguity set k has been the number one set from the LAMBDA function

Although not explicitly used in the minimum indicator equation, the ratio test described in Section 2.3.2 was used during the minimum indicator calculations. A candidate ambiguity set had to brake a minimum ratio specified in the initialization file before consideration as the minimum indicator. The ratio value was generally set at 1.5 or 2.0. As added protection against an incorrect set being selected as the minimum indicator, a candidate set must have been the top ambiguity set from LAMBDA at least once before it was considered as the minimum indicator. These last two checks further protected against a local minima set from being selected as the minimum indicator.

3.8 *Relative Position Update*

The relative position module ran at the raw measurement speed from the Javad® receivers. It would sit idle until a new set of measurements arrived. The precise centimeter-level navigation solution was the primary output from the algorithm. To provide this output at high data rates, an efficient method was needed for calculating the relative position. Much of the development for the relative position calculation has already been discussed in Section 3.7.1.

The relative position calculation was accomplished using only phase measurements since they were more precise and less noisy than the code measurements. The same phase correction from Equation (3.68) was applied to the incoming raw phase measurements.

Next, an ambiguity set was applied to the phase measurements as in Equation (3.70). The ambiguity set was either the current minimum indicator, or the blended ambiguity set from the MMAE. A flag in the initialization file determined which value was used as shown in Figure 3.5. Finally, the actual relative position calculation was Equation (3.72), repeated here for clarity:

$$\Delta \hat{\mathbf{x}} = \left[(\mathbf{H}^T \mathbf{H})^{-1} \mathbf{H}^T \right] \mathbf{z}_n \quad (3.84)$$

As shown in Figure 3.5, the “best” ambiguity set was only sent to the relative position module once a second. Therefore, the ambiguities could not change before the next floating filter update. The \mathbf{H} matrix in Equation (3.84) was also constant over the time period between floating filter updates. Since the unit line-of-sight vectors in the \mathbf{H} matrix changed very slowly, there was no degradation in the navigation solution by not calculating a new \mathbf{H} matrix each time new phase measurements arrived. This saved time and was much more efficient. So, the only change in Equation (3.84) from epoch-to-epoch was the incoming raw phase measurements, until a new filter update occurred. Once the floating filter update occurred, a new \mathbf{H} matrix and potentially a new set of ambiguities were sent to the relative position module. Potential problems such as cycle slips are discussed next.

3.9 Cycle Slips

The presence of cycle slips in the double-differenced phase observables posed a significant problem when trying to resolve the phase ambiguities. Newer receivers have done a remarkable job reducing cycle slips, and some even set a flag when a cycle slip occurs. The first step in fixing a cycle slip was accurate detection. After a cycle slip was detected, it had to be repaired or reset. The proper handling of cycle slips was imperative to achieving centimeter-level navigation accuracy.

3.9.1 Cycle Slip Detection. Since the algorithm ran in real-time, it needed to detect and handle any potential cycle slips properly. Originally, it was believed that the Javad® JNS100 had a cycle slip flag for the phase-lock loop (PLL). However, late in the testing phase it was discovered that the JNS100 was not only missing the cycle slip flag, but did not have any flags for the delay-lock loop (DLL) or PLL. The only functioning loop flag

was a loss-of-lock flag on the DLL. A combination of the loss-of-lock flag, triple-differenced phase observables, and a Doppler prediction method were used to detect cycle slips.

There has been much research in the area of cycle slip detection and correction [1–3, 14, 17]. However, most research revolved around dual-frequency receivers. Since the JNS100 was L1 only, the triple-differenced phase observables and Doppler prediction methods were the quickest, easiest way to detect cycle slips. The triple-differenced phase observables were created by:

$$\nabla\Delta\nabla\phi_{xy}^{ab} = \Delta\nabla\phi_{xy}^{ab}(t_i) - \Delta\nabla\phi_{xy}^{ab}(t_{i-1}) \quad (3.85)$$

where

$$\begin{aligned} \Delta\nabla\phi_{xy}^{ab} &= \text{Double-differenced phase measurement} \\ t_i &= \text{time at epoch } i \end{aligned}$$

After the triple-differenced observables were formed, one final difference was made for cycle slip detection. The difference between successive triple difference measurements was formed by:

$$\text{Error quantity} = \nabla\Delta\nabla\phi_{xy}^{ab}(t_i) - \nabla\Delta\nabla\phi_{xy}^{ab}(t_{i-1}) \quad (3.86)$$

where

$$\begin{aligned} \nabla\Delta\nabla\phi_{xy}^{ab} &= \text{Triple-differenced phase measurement} \\ t_i &= \text{time at epoch } i \end{aligned}$$

Expanding the triple-differenced phase observables from Equation (3.86) and collecting like terms, reveals why the triple-difference phase observables detect cycle slips. Expanding the triple-differenced phase term gave:

$$\begin{aligned} \nabla\Delta\nabla\phi_{xy}^{ab} &= \left[\frac{1}{\lambda} (\Delta\nabla r_{xy}^{ab} + \Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab}) + \Delta\nabla N_{xy}^{ab} \right] (t_i) - \\ &\quad \left[\frac{1}{\lambda} (\Delta\nabla r_{xy}^{ab} + \Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab}) + \Delta\nabla N_{xy}^{ab} \right] (t_{i-1}) \\ &= \frac{1}{\lambda} (\Delta\nabla r_{xy}^{ab} + \Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab}) (t_i) - \\ &\quad \frac{1}{\lambda} (\Delta\nabla r_{xy}^{ab} + \Delta\nabla T_{xy}^{ab} + \Delta\nabla I_{xy}^{ab} + \Delta\nabla m_{xy}^{ab} + \Delta\nabla v_{xy}^{ab}) (t_{i-1}) + \\ &\quad [\Delta\nabla N_{xy}^{ab}(t_i) - \Delta\nabla N_{xy}^{ab}(t_{i-1})] \end{aligned} \quad (3.87)$$

The the bracketed portion of the final line of Equation (3.87) holds the key. Since the ambiguities from one epoch to the next remained constant unless a cycle slip or loss-of-lock occurred, the bracketed term, $[\Delta\nabla N_{xy}^{ab}(t_i) - \Delta\nabla N_{xy}^{ab}(t_{i-1})]$, must be zero or a cycle slip had occurred. The difference between triple-difference ambiguities from adjacent epochs would also be zero unless a cycle slip occurred. Since all of the terms in Equation (3.87) canceled except the noise and range terms, the final difference canceled the range term (the noise terms could not be canceled). Therefore, if the error quantity in Equation (3.86) was much greater than zero, a cycle slip had occurred. Normally, if a cycle slip occurred, the difference from one epoch to the next would be an integer value equaling the size of the cycle slip. However, since all of the terms did not completely cancel, there was some inherent noise in the error quantity value.

The JNS100 exhibited some unusual behavior regarding cycle slips during the driving tests. Occasionally, a single-cycle cycle slip would occur on certain satellites. There did not appear to be any correlation between the satellite number, tracking loop or satellite elevation and the possibility of the strange cycle slip. However, there were tall buildings, trees and light poles in the driving area. The most likely cause was minor masking from one of the obstacles mentioned above. All cycle slips during flight testing were attributed to aircraft maneuvering.

Based on the driving tests, a 0.5 cycle threshold for triple-differenced observables was set. Even with the low threshold, very few cycle slips were detected during testing both on the ground and in flight. Ground testing revealed the triple-difference method was the best at detecting cycle slips. However, the Doppler prediction method (described below) was still used for redundancy.

The second technique used to catch cycle slips was the Doppler prediction method [47]. This approach generally works well when the data rate is relatively high. The smaller the data interval, the more accurate the Doppler measurements become. Since the algorithm ran at 20 Hz, the Doppler measurements from the Javad[®] receiver were a good prediction of the expected phase. The base equation relating carrier-phase to Doppler was:

$$\tilde{\phi}(t_i) = \phi(t_{i-1}) + (t_i - t_{i-1}) \Delta f(t_i) \quad (3.88)$$

where

$$\begin{aligned}\Delta f &= \text{Doppler measurement (Hz)} \\ \phi &= \text{Carrier-phase measurement (cycles)} \\ t_i &= \text{time at epoch } i \text{ (sec)}\end{aligned}$$

The predicted phase, $\tilde{\phi}(t_i)$, was then compared against the raw phase measurement from the same epoch. The difference between the two measurements was the prediction error.

$$\text{Error} = \tilde{\phi}(t_i) - \phi(t_i) \quad (3.89)$$

If the magnitude of the error from Equation (3.89) rose above a specified threshold, then a cycle slip had occurred. The algorithm in this thesis used a threshold of one cycle for the Doppler prediction method. This value was determined empirically.

The Javad® JNS100 periodically synchronized its internal clock to GPS time. This caused a spike in the Doppler measurement, which looked like a cycle slip. The problem was easily overcome since the clock shift was exactly one millisecond. The resulting Doppler change is:

$$\text{Doppler Shift} = \frac{c}{1000 \lambda} \quad (3.90)$$

where

$$\begin{aligned}c &= \text{Speed of light in vacuum (m/sec)} \\ \lambda &= \text{Carrier-phase wavelength (m/cycle, L1, L2, or Widelane)}\end{aligned}$$

Table 3.6 shows the values resulting from Equation (3.90).

Table 3.6: Doppler Change Caused by Clock Shift

Wavelength	Value
L1	1575420 cycles/sec
Widelane	347820 cycles/sec

If a loss-of-lock flag was set or either of the two previously mentioned cycle slip detection routines sensed a cycle slip, the associated satellite was flagged. The measurements and ambiguities from affected satellites were not used until the cycle slip was corrected. The next section describes how the algorithm handled cycle slip correction.

3.9.2 Cycle Slip Correction. No attempt was made to correct cycle slips in the algorithm. Instead, the affected ambiguity state was reset using Equation (3.11). Also, the associated covariance row and column in the \mathbf{P} matrix were zeroed with the ambiguity variance reset to the ambiguity value in Table 3.4. Before the reset occurred, the PLL and DLL were given ten seconds to stabilize. Since the flags for the PLL and DLL were turned off in the JNS100, this wait time was imperative to ensure accurate phase and code measurements. Testing showed significant improvement in the algorithm’s ability to handle cycle slips when the wait time was included.

The only exception to the above rule was when a cycle slip occurred on the reference satellite for the double-difference observables. If this happened, the entire filter was reset. Although there are better ways to handle such a situation (see [4] or Section 3.10 for an example), its occurrence was very low (the reference satellite was the highest elevated satellite, making cycle slips much less likely). A trade-off between added value and time required was made.

A cycle slip not only affects the floating filter, but also the MMAE, Minimum Indicator and relative position update. The next section discusses the impact of cycle slips on the MMAE and Minimum Indicator, followed by the relative position update.

3.9.3 MMAE and Minimum Indicator Cycle Slip Handling. Cycle slip correction was easy for the floating filter section, but much more involved and time consuming for the candidate ambiguity sets. If a cycle slip occurred, then every candidate ambiguity set was affected. The elemental filters from the MMAE and the Minimum Indicator section used the same candidate ambiguity sets for their calculations. Therefore, the candidate ambiguity sets were first corrected for cycle slips and then sent to each section.

First, the affected individual ambiguities were removed from each candidate set. After that, a temporary relative position calculation was made for each candidate set. The relative position was then used to estimate the affected ambiguities based on the raw phase measurements. In equation form, the process was similar to Section 3.7.1.

The first step was to remove the affected ambiguities from the candidate sets. The \mathbf{H} matrix and raw phase measurements were also modified since an ambiguity was removed.

The associated row from the \mathbf{H} matrix, and raw phase measurement were removed before the position calculation was performed. The estimated relative position conditioned on a particular ambiguity set was specified by:

$$\Delta \hat{\mathbf{x}}' = \left[((\mathbf{H}')^T \mathbf{H}')^{-1} (\mathbf{H}')^T \right] \mathbf{z}'_n \quad (3.91)$$

where

$$\begin{aligned} \mathbf{H}' &= \text{Modified } \mathbf{H} \text{ matrix with associated cycle slip rows removed} \\ \mathbf{z}'_n &= \text{Modified } \mathbf{z}_n \text{ vector with associated cycle slip raw phase} \\ &\quad \text{measurements and ambiguities removed} \end{aligned}$$

The temporary relative position, $\Delta \hat{\mathbf{x}}'$, was based only on the good (i.e., non-cycle slip) ambiguity values from the previous epoch. Now the \mathbf{H} matrix was returned to its original form. New phase measurements based on the estimated relative position and original \mathbf{H} matrix were calculated by:

$$(\Delta \nabla \phi)'' = \mathbf{H} \hat{\mathbf{x}}' \quad (3.92)$$

The estimated ambiguities were then calculated by subtracting the new phase measurements in Equation (3.92) from the corrected phase measurements in Equation (3.68):

$$\Delta \nabla \tilde{N} = (\Delta \nabla \phi)' - (\Delta \nabla \phi)'' \quad (3.93)$$

The ambiguities in Equation (3.93) would not necessarily be integers after the calculation (but they would be very close to integers). Therefore, the final step was to round each individual ambiguity term to an integer value.

One potential problem must be addressed, however. There was a possibility, when certain ambiguities were removed from the candidate ambiguity sets, that two or more sets would now be the same. If this happened, when the reverse transformation occurred, the ambiguity sets would still match. Since all of the candidate ambiguity sets must be unique, any duplicate sets must be removed. The ambiguity set with the lowest minimum indicator was kept, and all others were discarded.

3.9.4 Relative Position Update Cycle Slip Handling. Since the relative position module was designed for speed, there was no attempt to fix any detected cycle slips. If a cycle slip occurred, the associated satellite was flagged. The flag on that satellite remained until the next floating filter update and a reset was performed. No measurements from that satellite were used in the relative position calculation from the time the cycle slip was detected until it was reset.

If a cycle slip occurred on the reference satellite, the relative position update was stopped until the next floating filter update. This particular problem never occurred during testing. The only other time the relative position update was stopped was when there were fewer than four usable satellites, whether it was due to cycle slips or receiver antenna masking. Again, this never happened during testing.

3.10 Losing/Gaining Satellites

Another potential problem when running in real-time was the loss or gain of satellites. A satellite may be lost when it set below the elevation cutoff angle or was masked by an obstacle. Usually, satellites were only gained when they rose above the elevation cutoff angle. However, if the satellite was lost due to masking, then it could be reacquired when the obstruction was removed.

Losing a satellite was the easier of the two problems. If a satellite was lost, the associated ambiguity state and row and column from the \mathbf{P} matrix were removed from the floating filter. Also, the associated ambiguity state from the candidate ambiguities was removed. The candidate ambiguities were checked to ensure each was unique, as described in Section 3.9.3. These were the only changes required, aside from resizing other matrices such as the \mathbf{F} , \mathbf{R} , etc.

If a satellite was gained, the problem became more complex. The floating filter needed only to initialize a new ambiguity state in the state vector and insert rows and columns into the \mathbf{P} matrix. Next, the candidate ambiguities were corrected for the new satellite. The process was very similar to that of the cycle slip discussion in Section 3.9.3. But, the first step of removing a satellite was no longer needed. All of the other steps were identical. In

fact, a cycle slip was handled as if a satellite was lost and then gained on the same epoch. Lastly, the other affected matrices, \mathbf{F} , \mathbf{R} , etc., were resized.

3.11 Changing Reference Satellite

The possibility existed that the reference satellite would change during a data run. A quick transformation of the ambiguity states and associated covariances allowed the algorithm to retain all of the previous information when the change occurred. The ability to transform the state vector and covariance matrix was based on the formulation of the double-differenced phase observables. Within the double-differenced phase observables were the associated ambiguity terms. With a as the reference satellite, a particular ambiguity set for the phase observables was described by:

$$\begin{aligned}\Delta\nabla N_{xy}^{ab} &= \Delta N_{xy}^a - \Delta N_{xy}^b \\ \Delta\nabla N_{xy}^{ac} &= \Delta N_{xy}^a - \Delta N_{xy}^c \\ \Delta\nabla N_{xy}^{ad} &= \Delta N_{xy}^a - \Delta N_{xy}^d \\ \Delta\nabla N_{xy}^{ae} &= \Delta N_{xy}^a - \Delta N_{xy}^e\end{aligned}\tag{3.94}$$

Now if the reference satellite changed to e , the equations became:

$$\begin{aligned}\Delta\nabla N_{xy}^{ea} &= \Delta N_{xy}^e - \Delta N_{xy}^a \\ \Delta\nabla N_{xy}^{eb} &= \Delta N_{xy}^e - \Delta N_{xy}^b \\ \Delta\nabla N_{xy}^{ec} &= \Delta N_{xy}^e - \Delta N_{xy}^c \\ \Delta\nabla N_{xy}^{ed} &= \Delta N_{xy}^e - \Delta N_{xy}^d\end{aligned}\tag{3.95}$$

Modifying Equation (3.95) to:

$$\begin{aligned}\Delta N_{xy}^e &= \Delta N_{xy}^a + \Delta\nabla N_{xy}^{ea} \\ \Delta N_{xy}^b &= \Delta N_{xy}^e - \Delta\nabla N_{xy}^{eb} \\ \Delta N_{xy}^c &= \Delta N_{xy}^e - \Delta\nabla N_{xy}^{ec} \\ \Delta N_{xy}^d &= \Delta N_{xy}^e - \Delta\nabla N_{xy}^{ed}\end{aligned}\tag{3.96}$$

and substituting Equation (3.96) into Equation (3.94) gave:

$$\begin{aligned}
\Delta \nabla N_{xy}^{ab} &= \Delta N_{xy}^a - (\Delta N_{xy}^e - \Delta \nabla N_{xy}^{eb}) \\
\Delta \nabla N_{xy}^{ac} &= \Delta N_{xy}^a - (\Delta N_{xy}^e - \Delta \nabla N_{xy}^{ec}) \\
\Delta \nabla N_{xy}^{ad} &= \Delta N_{xy}^a - (\Delta N_{xy}^e - \Delta \nabla N_{xy}^{ed}) \\
\Delta \nabla N_{xy}^{ae} &= \Delta N_{xy}^a - (\Delta N_{xy}^e + \Delta \nabla N_{xy}^{ea})
\end{aligned} \tag{3.97}$$

Combining terms and reducing resulted in:

$$\begin{aligned}
\Delta \nabla N_{xy}^{ab} &= \Delta \nabla N_{xy}^{ae} + \Delta \nabla N_{xy}^{eb} \\
\Delta \nabla N_{xy}^{ac} &= \Delta \nabla N_{xy}^{ae} + \Delta \nabla N_{xy}^{ec} \\
\Delta \nabla N_{xy}^{ad} &= \Delta \nabla N_{xy}^{ae} + \Delta \nabla N_{xy}^{ed} \\
\Delta \nabla N_{xy}^{ae} &= -\Delta \nabla N_{xy}^{ea}
\end{aligned} \tag{3.98}$$

Finally, rearranging the equations so all of the new base satellite ambiguities are on the left gave:

$$\begin{aligned}
\Delta \nabla N_{xy}^{eb} &= \Delta \nabla N_{xy}^{ab} - \Delta \nabla N_{xy}^{ae} \\
\Delta \nabla N_{xy}^{ec} &= \Delta \nabla N_{xy}^{ac} - \Delta \nabla N_{xy}^{ae} \\
\Delta \nabla N_{xy}^{ed} &= \Delta \nabla N_{xy}^{ad} - \Delta \nabla N_{xy}^{ae} \\
\Delta \nabla N_{xy}^{ea} &= -\Delta \nabla N_{xy}^{ae}
\end{aligned} \tag{3.99}$$

The reason for the math may not be readily apparent at first, but all of the terms on the right hand side of Equation (3.99) are available from the first set of equations in Equation (3.94). Therefore, the pattern for changing the base satellite is now in a general formula form.

Using the example above and Equation (3.99), a transformation matrix was specified by:

$$\mathbf{W}' = \begin{bmatrix} 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{bmatrix} \tag{3.100}$$

Thus, to change the base satellite from a to e the following calculation was performed using the \mathbf{W}' matrix:

$$\begin{bmatrix} N_{xy}^{ea} \\ N_{xy}^{eb} \\ N_{xy}^{ec} \\ N_{xy}^{ed} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} N_{xy}^{ab} \\ N_{xy}^{ac} \\ N_{xy}^{ad} \\ N_{xy}^{ae} \end{bmatrix} \quad (3.101)$$

The \mathbf{W}' matrix would transform the ambiguities, but in order to transform the state vector and \mathbf{P} matrix, the relative position, velocity and acceleration states must be considered. When the reference satellite changes, there should be no change in any states except the ambiguities. Therefore, an identity matrix was combined with the \mathbf{W}' matrix to form the actual transformation matrix, \mathbf{W} . If the floating filter had the velocity and acceleration states, then using the previous example, the full \mathbf{W} matrix was:

$$\mathbf{W} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 & 0 \end{bmatrix} \quad (3.102)$$

The first nine rows and columns are the identity matrix, and the \mathbf{W}' matrix follows starting in $\mathbf{W}_{10,10}$. If the floating filter was in position-only mode, the identity matrix was only 3-by-3 followed by the \mathbf{W}' matrix in $\mathbf{W}_{4,4}$. Finally, the transformation of the floating filter

state vector, \mathbf{x} , and its associated covariance matrix, \mathbf{P} , was accomplished by:

$$\begin{aligned}\mathbf{x}' &= \mathbf{W} \mathbf{x} \\ \mathbf{P}' &= \mathbf{W} \mathbf{P} \mathbf{W}^T\end{aligned}\tag{3.103}$$

3.12 Summary

Chapter III has given a hardware design overview and discussed all of the major components of the software algorithm. First, the floating filter design and considerations were described, followed by the MMAE section. All of the MMAE structure and design was based on previous research by Henderson. The ambiguity set generation was performed by the LAMBDA method. Next, the new minimum indicator variable development was described. A quick and efficient method to calculate the relative position using only phase measurements was also described. Finally, other real-time aspects such as cycle slips, gaining/losing satellites and changing the reference satellite were explained. The next chapter will show the results from both ground and flight testing.

IV. Results and Analysis

4.1 Overview

This chapter presents the results and analysis from ground and flight testing of the algorithm developed in Chapter III. First, background information to include test methodology and data set descriptions are presented. Ground test results follow for a static test and a dynamic driving test. Next, flight test results from April and October 2005 are shown. Finally, special cases such as “pre-fit” and “post-fit” residual comparison, the floating filter in position-only mode, widelane measurements, and dynamic T_a and σ_a^2 are presented.

4.2 Test Background Information

4.2.1 Test Methodology. All laboratory and driving tests were conducted at Wright-Patterson Air Force Base (WPAFB) in Dayton, OH, from 5 January to 10 April 2005. The author used two personal vehicles to perform the driving tests, a 1994 Honda Accord and a 2004 Toyota Avalon. The test objectives during this phase, in order, were four-fold:

1. Verify the system would function in real-time.
2. Verify the data link was robust enough to handle minor obstructions and work at sufficient ranges (1000-3000 feet [305-914 m]).
3. Determine the algorithm’s ability to resolve ambiguities successfully in both static and dynamic situations.
4. Determine the best possible tuning parameters for the Kalman filter.

Based on the high cost of flight testing, a reliable system needed to be in place prior to the first flight test. Well over 100 ground tests were performed. Not all of the ground tests are discussed, only two to display the system’s performance under different vehicle dynamics.

All ground test objectives were met. In order to meet those objectives, the maneuvers listed in Table 4.1 were performed. The static sets were performed by

placing two antennas at a known relative position from each other, and then comparing the algorithm result to the true relative position. The true relative position for the static tests was measured by hand (using a ruler or tape measure), with an estimated accuracy of a few millimeters.

Table 4.1: Ground Maneuver Set for Relative GPS Position Testing

Maneuver	Nominal Conditions	Remarks
Static Tests	1 to 20 m	True relative position hand measured
Driving Tests	3 to 1000 m, 0 to 55 mph	True relative position from GrafMov

Note: m = meters and mph = miles-per-hour

For the dynamic driving tests, the lead vehicle maintained a constant speed, depending on the speed limit, while the trail vehicle varied speed to close and open the distance between the cars. Occasionally, the lead vehicle stopped for traffic lights, and the trail vehicle simulated a “rejoin”. The truth data for the driving tests came from the same GrafMov software described in Section 3.2.

All flight testing was conducted on either a C-12C or Calspan Learjet 24 at Edwards Air Force Base (EAFB), CA by the Air Force Test Pilot School (TPS). Students at TPS ran a Test Management Project (TMP) as their graduation exercise. Two separate TMPs supported this research. The first project, designated “Lost Wingman”, flew in April 2005 [16]. The second project, designated “No Gyro”, flew in October 2005 [50].

Although not essential to the work in this thesis, the MT9 microelectromechanical systems (MEMS) inertial measurement unit (IMU) was an integral part of the final system used in October 2005. Little testing was done on the MT9 filter algorithm in the lead aircraft prior to first flight. Therefore, many flight maneuvers were included to test the IMU operation during the “Lost Wingman” TMP. The test maneuvers included items such as, bank-to-bank rolls, straight-and-level unaccelerated flight (SLUF), steady heading side slips (SHSS), level accelerations / decelerations, climbs / descents, and discrete heading changes. The inertial information was critical

for the success of the “No Gyro” project. Based on the results from the “Lost Wingman” TMP, the MT9 was replaced by a similar MEMS IMU developed by Microbotics, Inc., the MIDG II INS/GPS, for use in the October flight tests [40].

The results of the DGPS algorithm from both TMPs are discussed later. However, the MT9 and MIDG II operations are outside the scope of this thesis, and no IMU results are presented. A detailed discussion of the MT9 results are in [16]. Also, further information on the MT9, with similar flight testing, can be found in [22]. The MIDG II results are in [50]. The flight tests had similar objectives to the ground testing. The general objectives for the flight tests were:

1. Determine relative position accuracy.
2. Determine data link robustness.
3. Determine IMU accuracy.

All flight test objectives were met. The maneuver set used for flight testing is shown in Table 4.2. The aircraft configuration for all test points was gear up and flaps up. In addition to the maneuvers listed in Table 4.2, the MT9 maneuvers described in the paragraph above were flown between 120 to 230 knots indicated airspeed (KIAS), and 8,000 to 20,000 feet pressure altitude (PA).

Table 4.2: C-12C Aircraft Maneuver Set for Relative GPS Position Testing

Maneuver	Nominal Conditions	Remarks
Pre-contact*	190 KIAS, 10,000 ft PA	Tolerance: ± 5 kts, ± 100 ft
Contact*	190 KIAS, 10,000 ft PA	Tolerance: ± 5 kts, ± 100 ft
Observation*	190 KIAS, 10,000 ft PA	Tolerance: ± 5 kts, ± 100 ft
Pre-contact to Contact	190 KIAS, 10,000 ft PA	Tolerance: ± 5 kts, ± 100 ft
Observation to Pre-contact	190 KIAS, 10,000 ft PA	Tolerance: ± 5 kts, ± 100 ft
Pre-contact to Observation	190 KIAS, 10,000 ft PA	Tolerance: ± 5 kts, ± 100 ft

*Note: Stabilized maneuvers

4.2.2 Data Sets. The data sets are broken down into ground, flight, and special sections. In the ground section, each data run was one complete set. There are

one static case and one dynamic case presented. In the flight section, the data sets were one hour long to reduce file size and to test the filter initialization/convergence during flight conditions. Generally, there were two data sets for each flight. Finally, the special section discusses areas such as, “pre-fit” and “post-fit” residual comparison, the floating filter in position-only mode, widelane measurements, and dynamic T_a and σ_a^2 . For all test cases, the error was defined as:

$$\hat{x}_{\text{error}} = \hat{x}_{\text{type}} - \hat{x}_{\text{true}} \quad (4.1)$$

where

$$\begin{aligned} \hat{x}_{\text{type}} &= \text{estimated relative position from the floating filter } (\hat{x}_{\text{float}}), \\ &\quad \text{MMAE } (\hat{x}_{\text{MMAE}}), \text{ or minimum indicator } (\hat{x}_{\text{Min Ind}}) \\ \hat{x}_{\text{true}} &= \text{true relative position from GrafMov software, or} \\ &\quad \text{hand measured} \end{aligned}$$

Table 4.3 summarizes the details of each data set. Not all of the data sets listed in Table 4.3 are discussed in this chapter. Only those sets where noteworthy events occurred are mentioned. Every data set, with all of their associated plots, are in Appendix A.

4.2.3 Data Reduction. As a reminder, the truth data for the relative position and double-difference ambiguities were calculated using the GrafMov 7.50 software, developed by Waypoint Consulting Inc., as described in Section 3.2. The GrafMov program used the same measurement observables from the Javad® JSN100 receivers that were used during the real-time execution. It would have been preferable to have independent GPS measurements for use in the GrafMov software as the truth data. Although the GAINR system was collecting GPS measurement data on the C-12Cs during the flight tests, the data format was proprietary and the author did not have access to the raw measurement data.

Table 4.3: Data Set Summary

Data Set	Length (min:sec)	Date	Time (local)	Location	Description
<i>G1</i>	59:57	25 March 05	1731	WPAFB	Static Ground Test
<i>G2</i>	59:54	30 March 05	1547	WPAFB	Dynamic Ground Test
<i>F1.1</i>	59:32	18 April 05	0913	EAFB	Flight Test 1.1
<i>F1.2</i>	59:32	18 April 05	1021	EAFB	Flight Test 1.2
<i>F2.1</i>	42:54	27 April 05	0942	EAFB	Flight Test 2.1
<i>F2.2</i>	23:46	27 April 05	1103	EAFB	Flight Test 2.2
<i>F3.1</i>	59:33	5 October 05	0920	EAFB	Flight Test 3.1
<i>F3.2</i>	37:00	5 October 05	1020	EAFB	Flight Test 3.2
<i>F4.1</i>	38:08	6 October 05	1112	EAFB	Flight Test 4.1
<i>F5.1</i>	59:33	11 October 05	1516	EAFB	Flight Test 5.1
<i>F5.2</i>	33:47	11 October 05	1617	EAFB	Flight Test 5.2
<i>F6.1</i>	59:33	12 October 05	0918	EAFB	Flight Test 6.1
<i>F6.2</i>	51:50	12 October 05	1019	EAFB	Flight Test 6.2
<i>F6.3</i>	59:33	12 October 05	1440	EAFB	Flight Test 6.3
<i>F6.4</i>	30:55	12 October 05	1540	EAFB	Flight Test 6.4
<i>F7.1</i>	59:33	13 October 05	1007	EAFB	Flight Test 7.1
<i>F7.2</i>	41:34	13 October 05	1107	EAFB	Flight Test 7.2
<i>F8.1</i>	59:33	14 October 05	0901	EAFB	Flight Test 8.1
<i>F8.2</i>	39:28	14 October 05	1001	EAFB	Flight Test 8.2

The GrafMov software was capable of running the data sets both forward and backward to get the best solution. After processing, the forward and backward solutions were combined using variance weighting. The combined solution was saved in a *.cmb file with only the epochs based on fixed integers. Usually, only a few epochs, out of approximately 72,000 epochs for an hour-long data run, were skipped by the GrafMov software since it could not resolve the ambiguities. The author developed a MATLAB[®] routine to read the *.cmb file and extract the precise relative position, relative velocity, and relative acceleration for each epoch. The MATLAB[®] routine then created three other files which had a GPS time stamp combined with the precise relative position, relative velocity, and relative acceleration. The final step was matching the GPS times for each epoch between the test data and the truth data. If there was

an epoch missing in the truth data, the epoch was skipped in the test data until the next valid truth epoch was available.

The precise relative position, in ECEF coordinates, was on the line designated “Ecf:” in the *.cmb file. The precise relative velocity and relative acceleration were in the local level frame. The precise relative velocity was the second set of three values on the “Loc:” line, and the precise relative acceleration was on the “Acc:” line (both in the order East, North, Up).

In addition to the *.cmb file, the author also used the *.fwd and *.rev files to get the ambiguities at each epoch. The ambiguities were extracted manually from each file. The GrafMov software listed the ambiguities in a normalized form in the output files. The output files had a “Bls:” section in each epoch that listed the number of visible satellites, the base satellite for the double-difference calculations, and each satellite with its normalized double-difference ambiguity. The author contacted Waypoint Consulting Inc. to determine the formula to extract the true double-difference (DD) ambiguities, and was given the following:

$$\text{DD Amb}_{\text{true}} = (\text{Sat \#} - \text{Base Sat \#})1000 - \text{DD Amb}_{\text{normalized}} \quad (4.2)$$

where

$$\begin{aligned} \text{DD Amb} &= \text{double-difference ambiguity} \\ \text{Sat \#} &= \text{Satellite Number (e.g. 1-32)} \end{aligned}$$

Next, ground test results will be discussed.

4.3 Ground Tests

Two ground tests are discussed in this section. The first was a static case taken at the AFIT navigation laboratory on 25 March 2005. Next, a dynamic driving case is shown which shows some of the errors discussed in Chapter III.

4.3.1 Case G1: Static Ground Test. Case G1 is the only static data set discussed, other than the case involving widelane observables. The algorithm easily handled the static case. Figure 4.1 displays the floating filter’s position error along with the floating filter’s 1σ estimate. The dashed lines represent the floating filter’s 1σ estimate.

The title of each subplot contains the associated statistics for that plot. Additionally, the distance root mean square (DRMS) and mean radial spherical error (MRSE) are listed in the first (East) subgraph title. The DRMS can be thought of as a horizontal 2-dimensional indicator of error, and the MRSE as a spherical 3-dimensional indicator of error. These two values define system performance better than the mean and standard deviation alone. The MRSE was used to determine the algorithm accuracy for each data run. This value was then compared against the centimeter-level requirement.

After an initial convergence period of approximately 10 minutes, the floating filter “found” the correct relative position. As expected, the floating filter was very accurate for a static case with an extremely short baseline (1 meter). Based on the floating filter’s 1σ estimation, the floating filter required more tuning for this particular case. However, the floating filter was designed to work with a variety of vehicle dynamics and therefore was not perfectly tuned for the static case. The floating filter was tuned to give the best possible outcome from the minimum indicator and the best MMAE weighted position estimate. In order to do this, the LAMBDA algorithm must return the true ambiguity set. The floating filter was only a tool to help the LAMBDA function. Ensuring the LAMBDA algorithm returned the true ambiguity set in the quickest possible time was the top priority, even at the expense of not tuning the floating filter for a theoretically accurate covariance matrix.

Next, a single satellite’s measurement data is shown in Figure 4.2. The other satellites’ plots are very similar and can be found in Appendix A. A small amount of code multipath can be seen in the cyclic nature of the pseudorange residuals. Also,

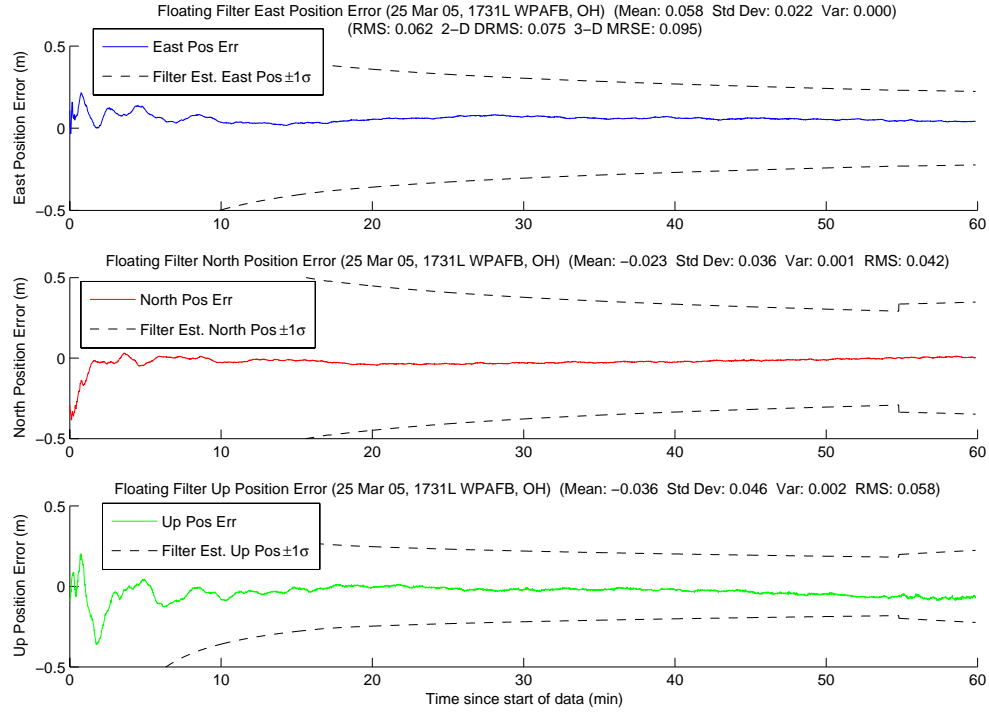


Figure 4.1: Case G1: Floating Filter Position Errors

the decrease in the 1σ for the ambiguity estimate at 18.5 minutes was due to a change in the reference satellite from pseudo-random noise (PRN) 18 to PRN 21. Figure 4.3 shows the number of visible satellites throughout the data run.

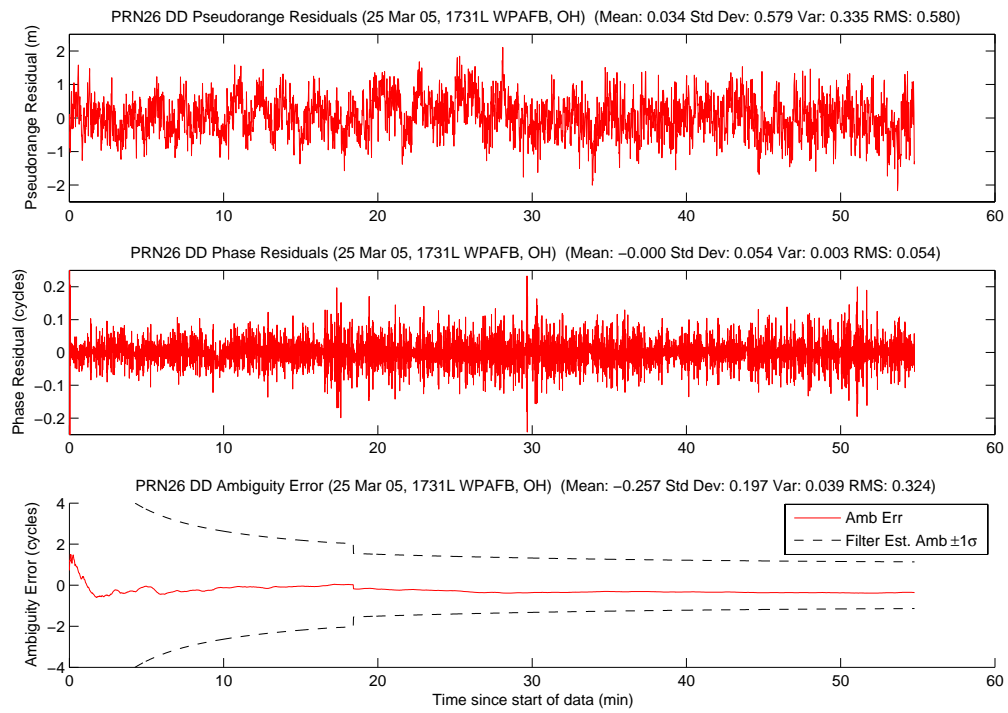


Figure 4.2: Case G1: Satellite 26 Measurements

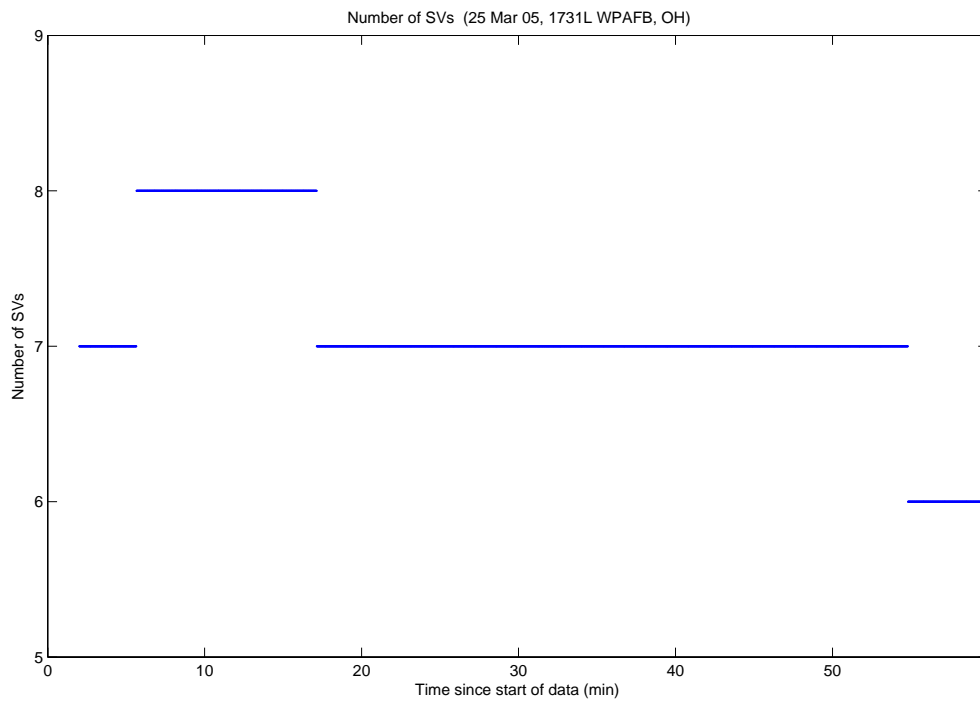


Figure 4.3: Case G1: Visible Satellites

After examining the output files from the data run, there were a total of 21 unique ambiguity sets returned from the LAMBDA algorithm. Since there was a two minute convergence time for the floating filter prior to sending the floating filter ambiguity estimates and their associated covariance to the LAMBDA algorithm, and a two minute wait time prior to performing the minimum indicator calculations, the earliest the correct ambiguity set could be discovered was four minutes into a data run. In the first pair of candidates returned from LAMBDA, the true set was the second candidate set. The top candidate was later culled from the ambiguity bank as the floating filter converged to the correct relative position. This was the only candidate pruned during the data run. The minimum indicator initially selected the wrong ambiguity set at four minutes. However, at the next least squares fit epoch, ten second later, the minimum indicator correctly selected the true ambiguity set. The minimum indicator did not change throughout the rest of the data run. Also, the true ambiguity set was returned from the LAMBDA function as the top set, 3411 out of 3600 epochs (94.8 percent). Finally, the highest ratio reached by the true ambiguity set was 31.8.

All the candidate ambiguity set's SOS residuals are shown in Figure 4.4 and the probabilities in Figure 4.5. The EF[#] in the legend of the probability plot shows which elemental filter (EF) was associated with the ambiguity set displayed beside it. A slash in the EF number section indicates a lower EF candidate set was pruned and the elemental filter's new number is shown. In this example, the number one EF was pruned. Therefore, all the elemental filters moved up one position. The parenthesis in the legend are the PRN numbers with their associated double-difference ambiguity. The true ambiguity set is explicitly identified with an arrow when required. To reduce clutter, only candidate sets which rose above 49 percent are displayed.

One interesting aspect of Figure 4.5 is when the MMAE "found" the true ambiguity set. This occurred at four minutes, approximately the same time the minimum indicator variable determined the correct set. Since each technique was independent of the other, this did not always occur.

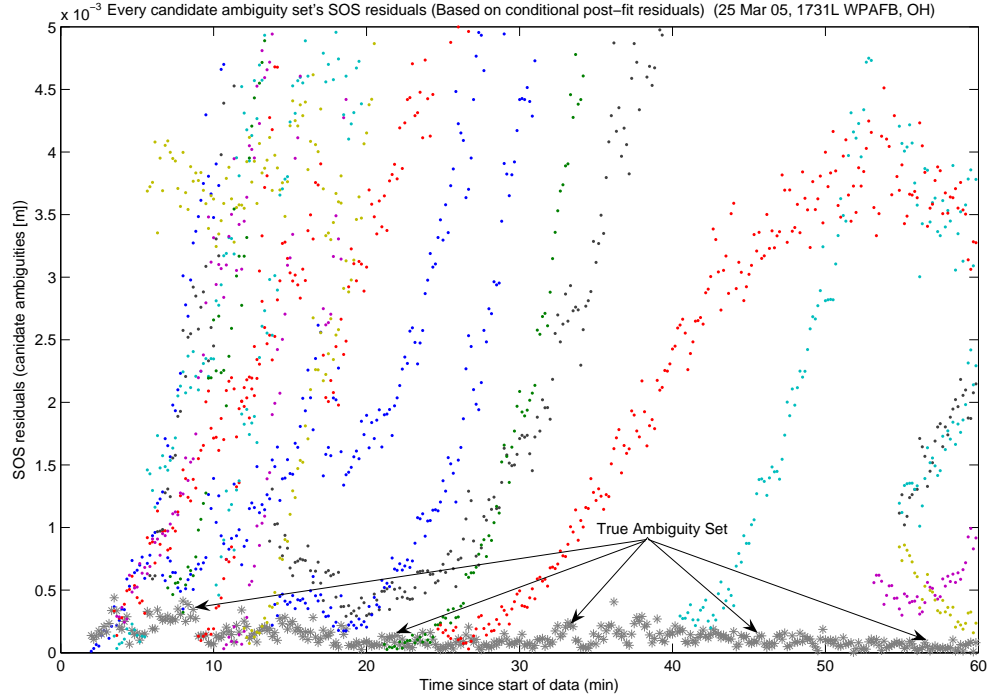


Figure 4.4: Case G1: SOS Residuals

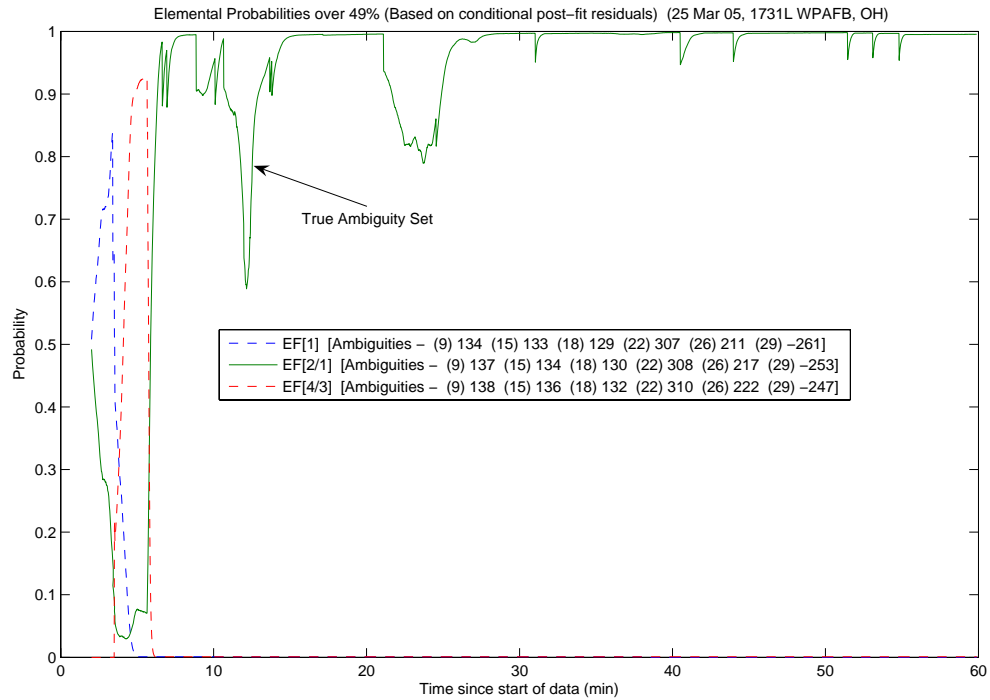


Figure 4.5: Case G1: Conditional Probabilities for Selected MMAE Elemental Filters

An examination of the SOS residuals and conditional probability plots reveal more interesting information. To help identify the true ambiguity set, its SOS residuals are denoted by grey asterisks in all the plots. By 10 minutes into the data run, the true set became apparent to the human eye. As the floating filter continued to gather more information from the measurements, and the covariance matrix “shrunk”, the true set stood out even more, especially after 25 minutes. However, some erroneous sets still managed to approach the true ambiguity set at 42 and 58 minutes in the run. As the satellite geometry changed, certain erroneous ambiguity sets’ measurements appeared to “fit” the position solution better due to errors such as multipath and noise. This reduced their SOS residuals to a level equal to, or lower than, that of the true set.

Two important times to note on the residual plot occur at 11 and 22 minutes. At both locations there are erroneous candidate sets with residuals that are lower than the true ambiguity set. In each case, the SOS residuals immediately begin to grow and rise above those of the true set. These locations correspond to significant dips in the true ambiguity set’s conditional probability, as shown in Figure 4.5. There are smaller probability dips which correspond to other wrong candidate sets’ SOS residuals approaching those of the true set. A recommendation for improvement of the MMAE position estimation is discussed in Chapter V.

Dips in the true ambiguity set’s probability mean the MMAE blended solution included an incorrect ambiguity set. Since the minimum indicator correctly chose the true set at 4.166 minutes and never lost “lock”, the difference between the MMAE and minimum indicator relative position estimates shows when the blended solution of the MMAE is “pulled” away from the true position. Figure 4.6 shows the difference between the MMAE and the minimum indicator position solutions. The figure starts at 4.5 minutes; directly after the true ambiguity set was found by the minimum indicator variable. Any difference between the two positions was due to an incorrect ambiguity set blending with the true ambiguity set. The effect of the reduced proba-

bility is visible at 11 and 22 minutes. Also, the minor dips in the probability can be seen in the position error later in the run.

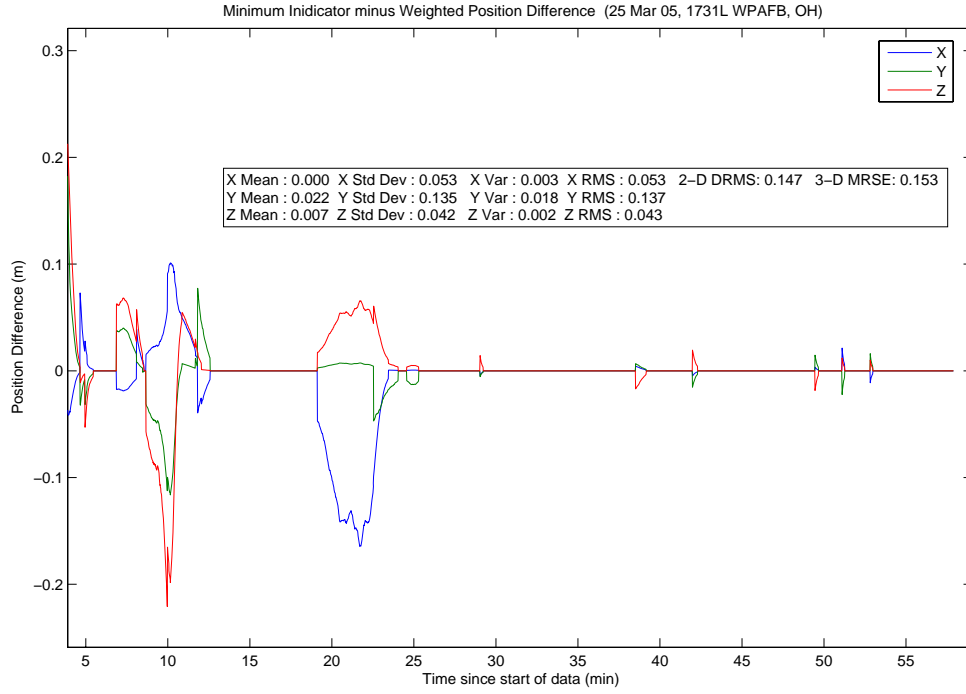


Figure 4.6: Case G1: Minimum Indicator and MMAE Position Difference

Finally, Figure 4.7 displays the MMAE weighted position error. The MMAE position solution was “pulled” off at 10 and 22 minutes, corresponding to the probability dips in Figure 4.5. Figure 4.8 shows the relative position error for the minimum indicator variable. The minimum indicator performed as expected since it was “locked” to the true ambiguity set during the entire data run. The scale on the minimum indicator error plot is one order of magnitude less than on the MMAE error plot. Lastly, Table 4.4 summarizes the floating point, MMAE and minimum indicator mean position error, standard deviation, 2-D DRMS and 3-D MRSE.

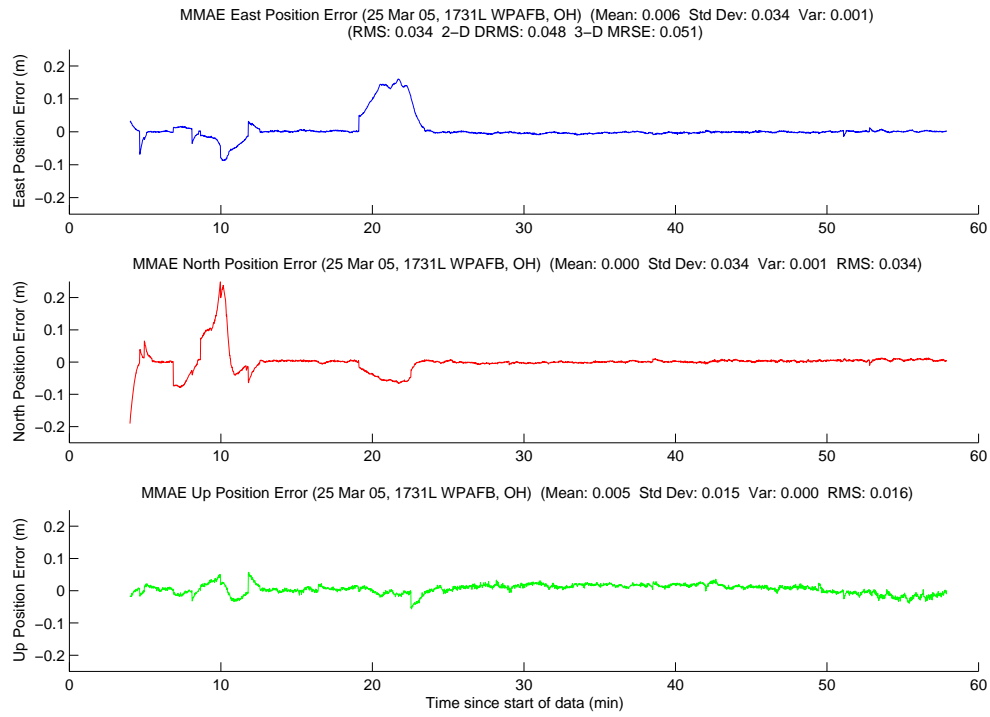


Figure 4.7: Case G1: MMAE Position Error

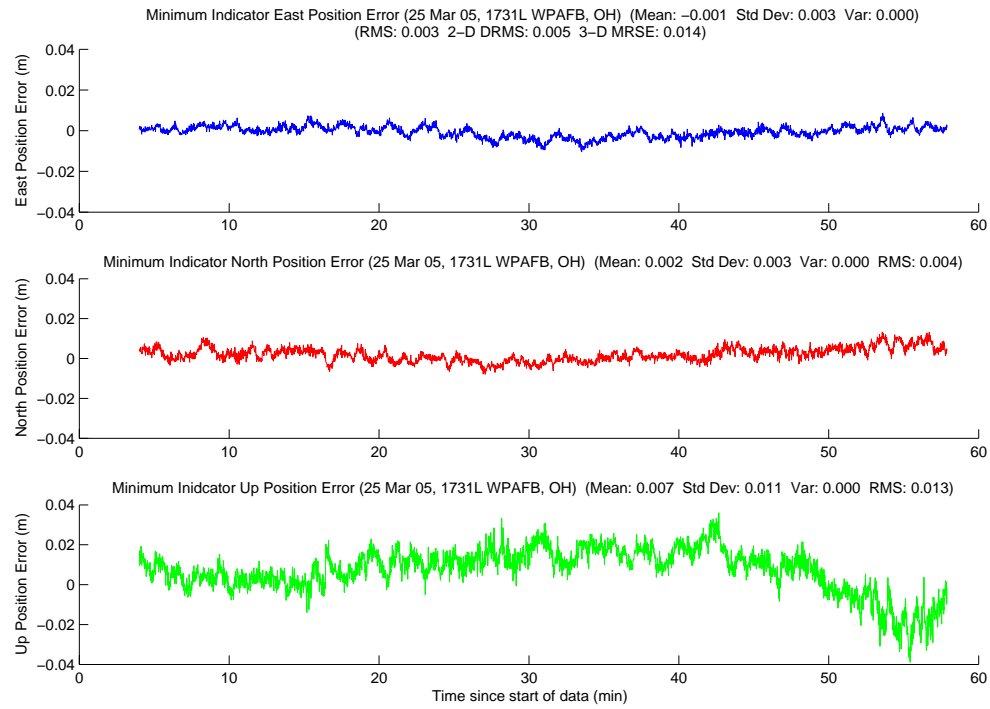


Figure 4.8: Case G1: Minimum Indicator Position Error

Table 4.4: Case G1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.058	0.022	-0.023	0.036	-0.036	0.046	0.075	0.095
MMAE	0.006	0.034	0.000	0.034	0.005	0.015	0.048	0.051
Min. Ind.	0.001	0.003	0.002	0.003	0.007	0.011	0.005	0.014

4.3.2 Case G2: Dynamic Driving Test. The second ground test was a dynamic driving test conducted on 30 March 2005. During the test, the cars ranged from 3.5 to 77.6 meters of separation. This test shows one of the single-cycle slips mentioned in Chapter III, and its effect on the floating filter, MMAE, and minimum indicator. The minimum indicator selected the true ambiguity set immediately at four minutes. There were no other sets chosen before the next true ambiguity set was discovered at 50.8 minutes into the data run. The true ambiguity sets were the only sets selected by the minimum indicator variable.

Throughout the data run, there were a total of 39 unique ambiguity sets returned from the LAMBDA algorithm. The first true ambiguity set was the top candidate set returned from the LAMBDA algorithm on the first call. Later in the run, LAMBDA returned the second true ambiguity set 30 seconds after the cycle slip was detected. In all, the true ambiguity set was returned as the top set from LAMBDA, 2984 of 3474 epochs (85.9 percent), and reached a maximum ratio of 4.4.

The floating filter performance is shown in Figure 4.9. The floating filter had trouble converging initially. The most likely cause was the lower number of visible satellites when the filter started. The filter initialized with seven satellites, but then lost a satellite three minutes later. Six satellites were about the minimum required for the floating filter to operate efficiently with only L1 observables. At 18 minutes into the data run, the filter began converging properly. This coincides with two satellites rising into view within three minutes of each other; denoting the importance of the number of visible satellites and their associated geometry.

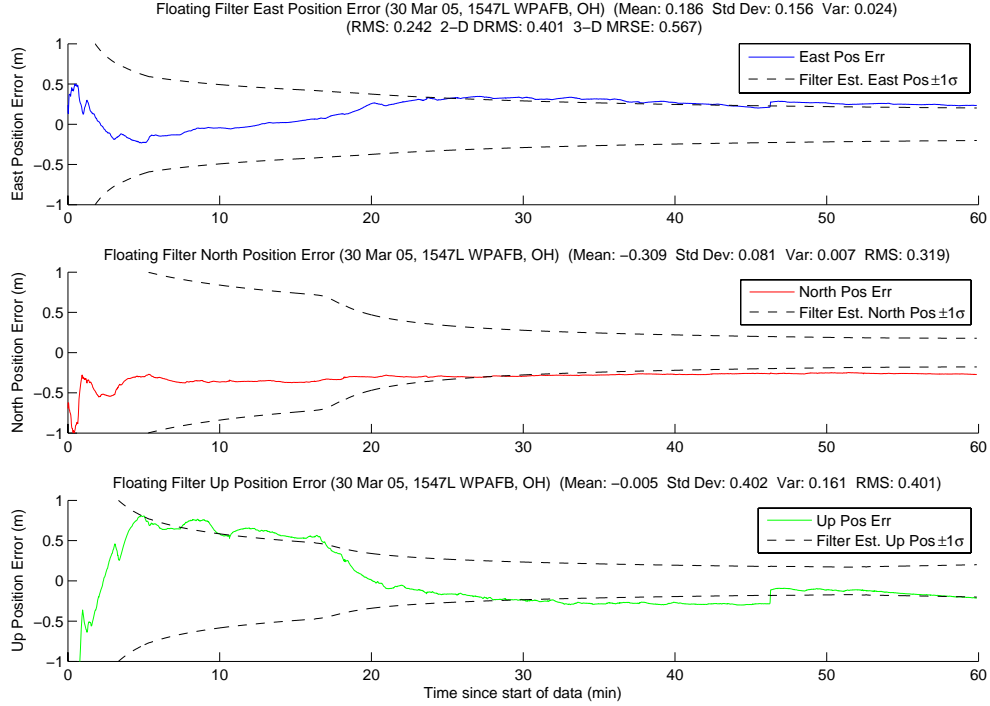


Figure 4.9: Case G2: Floating Filter Position Errors

Figure 4.10 displays PRN 6 measurement information and the floating filter's ambiguity estimate. The cycle slip that caused the minimum indicator problems is visible at approximately 46 minutes with a “jump” in the floating filter ambiguity error. The floating filter's 1σ estimate of the ambiguity state converged in one epoch. Thus, there was only a tiny bump in the $\mathbf{P}_k(t_i^+)$ 1σ value at the cycle slip, even though PRN 6 variance was reset. Finally, the phase residuals were significantly higher than those from Case G1, in Figure 4.2 (note scale difference). The higher phase residuals were due to the vehicle dynamics.

Figure 4.11 shows the number of visible satellites. The satellite dropouts were due to cycle slips (real or erroneous) or loss-of-lock. The only satellite vehicle (SV) dropouts due to actual cycle slips were at approximately 46 and 51 minutes. PRN 6 had a cycle slip at 46 minutes, which initially caused the minimum indicator to be incorrect. A few minutes later, PRN 22 also had a cycle slip, but was reset before the minimum indicator determined the true ambiguity set.

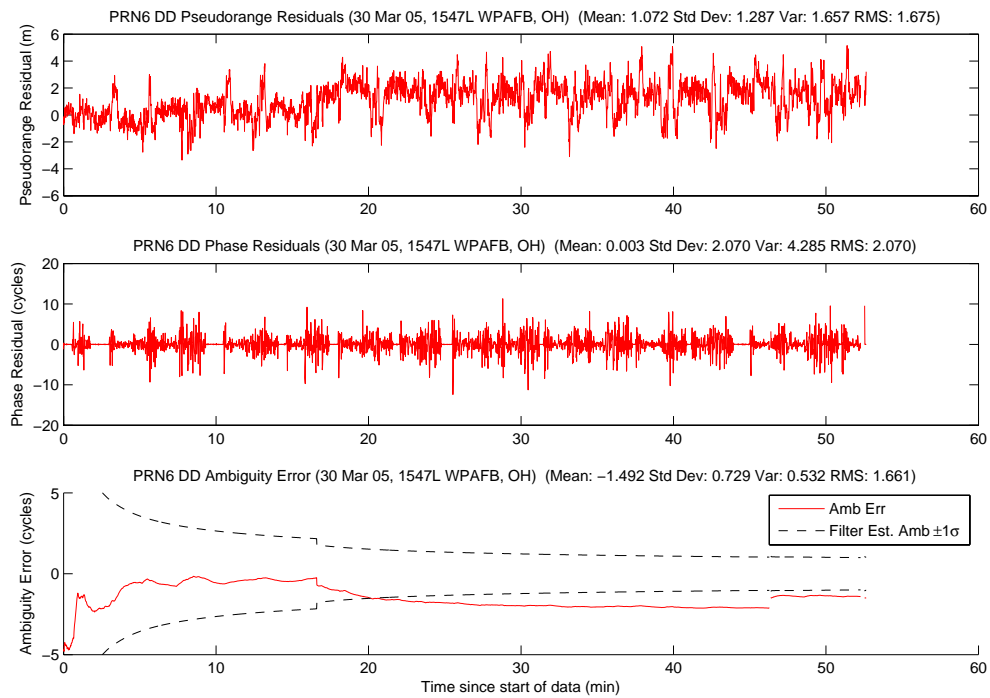


Figure 4.10: Case G2: Satellite 6 Measurements

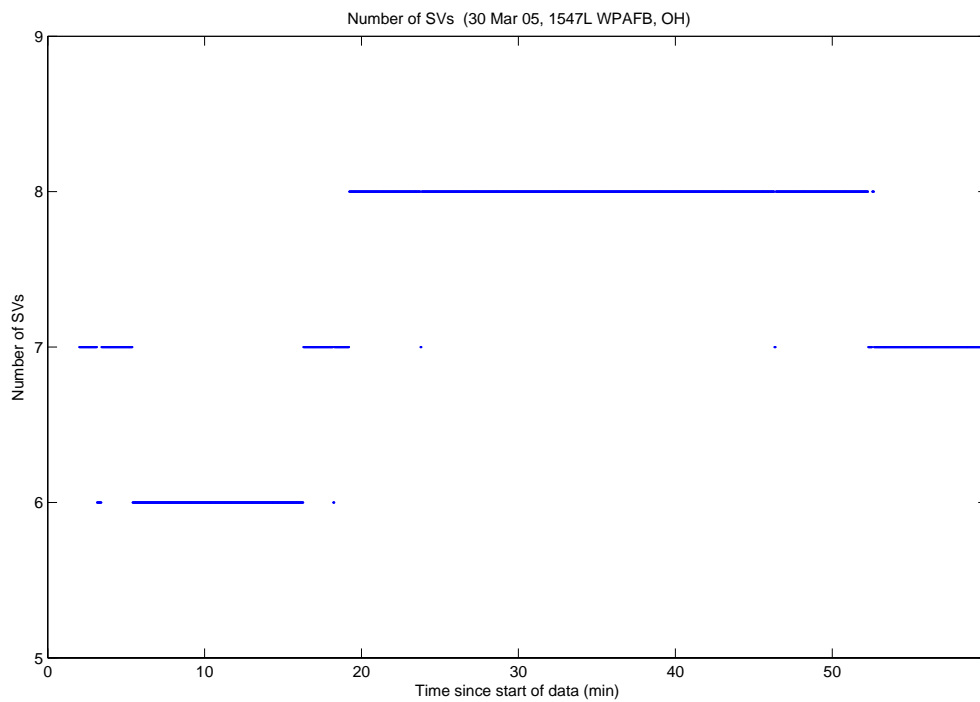


Figure 4.11: Case G2: Visible Satellites

The residual plot, Figure 4.12, is expanded to show the true ambiguity set's residuals better. The probability plot, Figure 4.13, has three points of interest. The first is at approximately three minutes where the MMAE determined the true ambiguity set. This was a full minute before the minimum indicator. Then at 10 minutes, another false candidate set absorbed 50 percent of the conditional probability, which corresponds to a rise in the true set's residuals in Figure 4.12. Finally, at 46.33 minutes, the true ambiguity set's probability immediately dropped when the cycle slip occurred. After the cycle slip, numerous erroneous ambiguity sets vied for the released probability. After 30 seconds, the second true ambiguity set was returned from LAMBDA, and absorbed all of the probability a minute later as seen in Figure 4.14.

The minimum indicator did not select the second correct set until 50.8 minutes into the run, about four minutes after the cycle slip. The 30 seconds for true ambiguity set to be returned from LAMBDA, coupled with the initial 2 minute SOS residuals acquisition time, means it took the minimum indicator 1.5 minutes to find the true set. The 1.5 minute delay was primarily driven by the fact that the minimum indicator divides by the summation of total number of times a candidate was returned as the top set from LAMBDA and the number of epochs it has been "alive". Since the first true ambiguity set had been the top set over 2200 epochs and was "alive" for approximately 3200 epochs, it took time for its minimum indicator value to grow (based on its much higher residuals) above the new true ambiguity set. Time constraints limited the author's ability to address this problem in the code, but a potential fix is discussed in Chapter V.

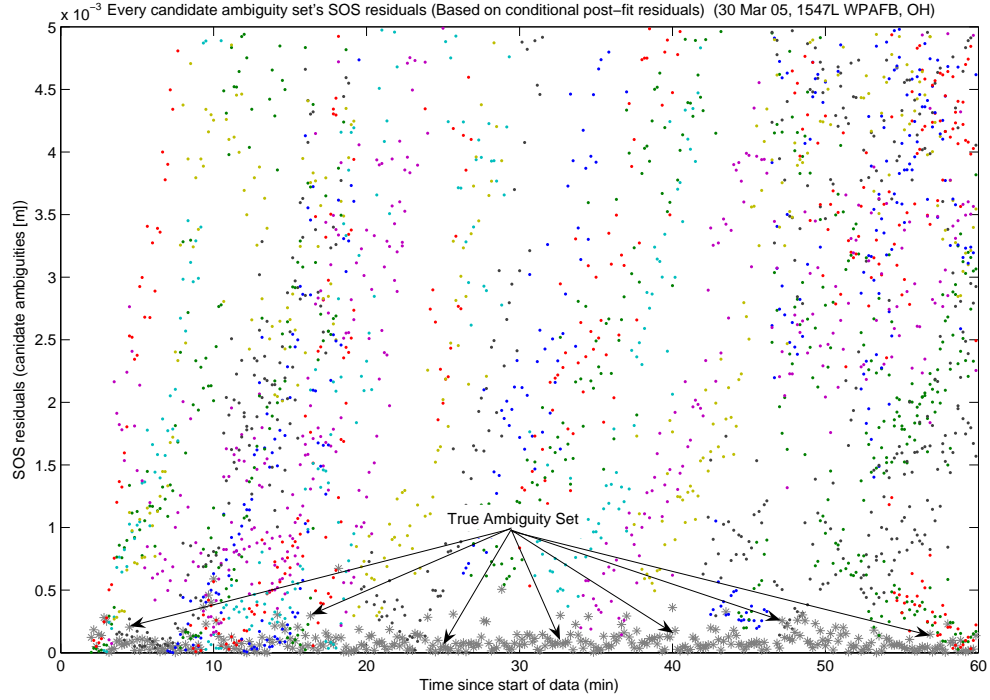


Figure 4.12: Case G2: SOS Residuals

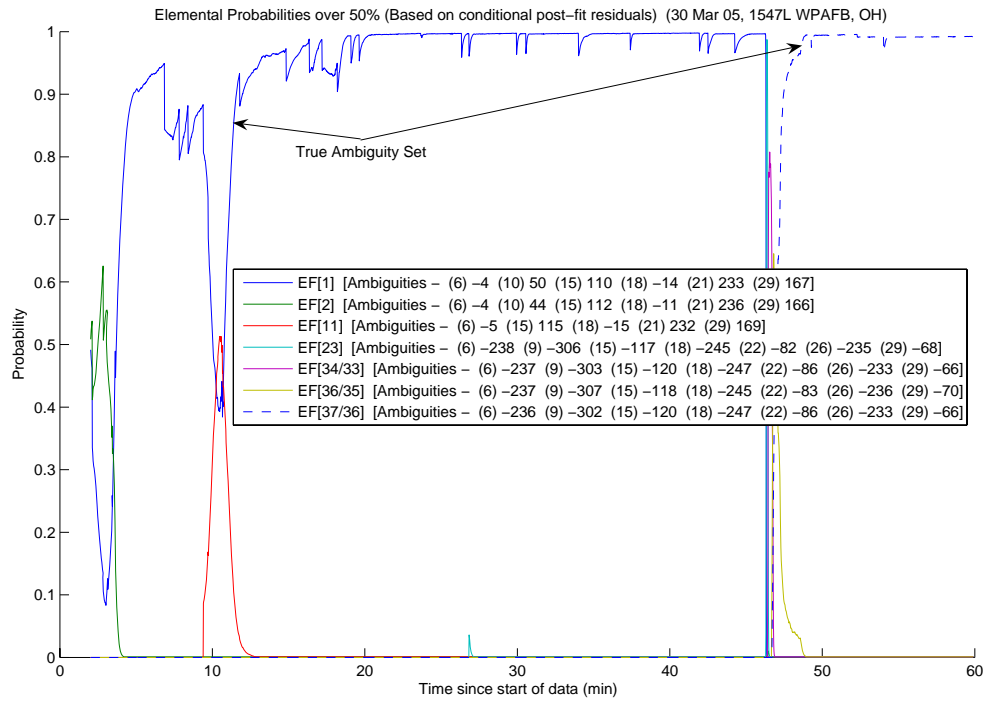


Figure 4.13: Case G2: Conditional Probabilities for Selected MMAE Elemental Filters

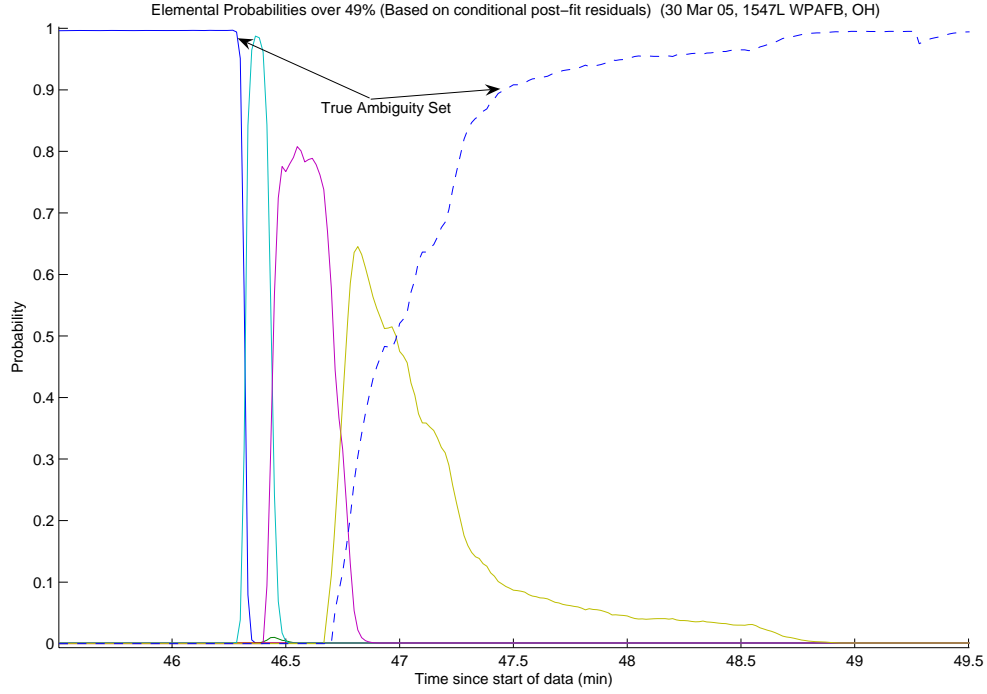


Figure 4.14: Case G2: Conditional Probabilities for Selected MMAE Elemental Filters Expanded

Figure 4.15 shows the MMAE position error. Since the MMAE quickly determined the true ambiguity set, it is relatively accurate from the beginning of the data run, unlike the other ground case. However, the probability dip at 10 minutes and the cycle slip at 46 minutes can both be seen in the plot.

Finally, Figure 4.16 displays the position error plot for the minimum indicator variable and Table 4.5 gives a summary of each method's position error. Even with the cycle slip, the minimum indicator remained reasonably accurate. Since the filter was using L1 observables, and only one satellite had a single-cycle cycle slip, the relative position error remained small. The magnitude of the error was approximately 18.7 centimeters. As expected, the error was close to the L1 wavelength of 19.0 centimeters. The reduced error from cycle slips and slightly better overall accuracy are the main advantages of L1 observables over widelane observables.

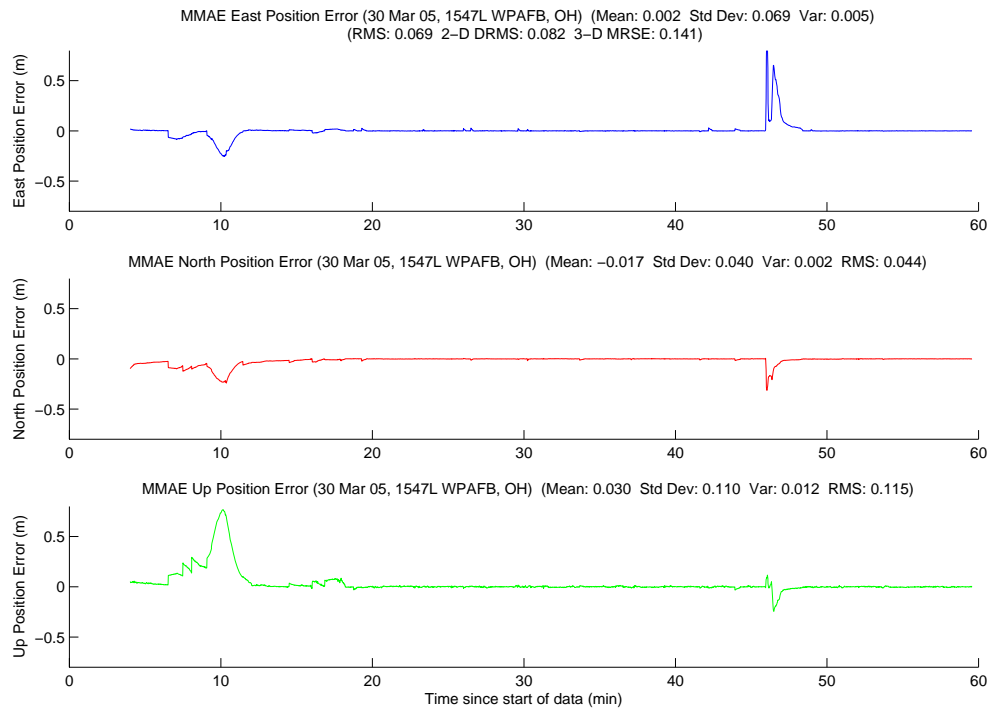


Figure 4.15: Case G2: MMAE Position Error

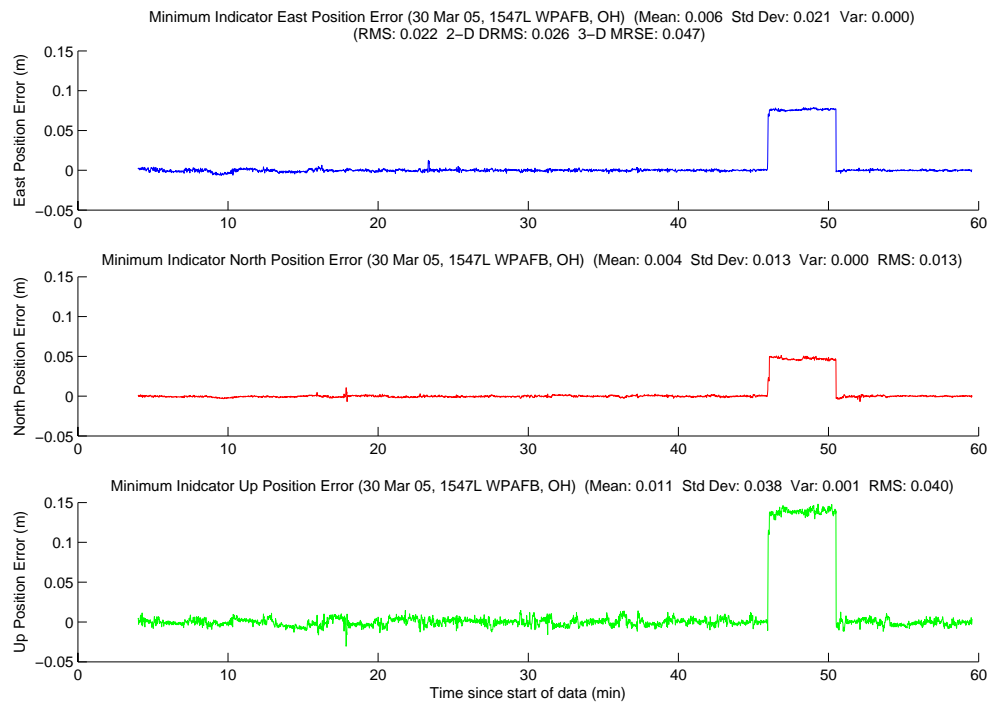


Figure 4.16: Case G2: Minimum Indicator Position Error

Table 4.5: Case G2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.186	0.156	-0.309	0.081	-0.005	0.402	0.401	0.567
MMAE	0.002	0.069	-0.017	0.040	0.030	0.110	0.082	0.141
Min. Ind.	0.006	0.021	0.004	0.013	0.011	0.038	0.026	0.047

4.3.3 Ground Test Summary. Two ground tests were discussed with a total run time of approximately two hours. During the two hours there were five SV dropouts, all during the dynamic driving test. Two were due to actual cycle slips and three were due to erroneously detected cycle slips. The algorithm correctly handled all of the SV dropouts. The base satellite changed twice during the two runs. Once more, the algorithm correctly handled the base satellite changes. Most importantly, all ground objectives were met. The system operated in real-time, the data link was verified to useful ranges, the algorithm successfully resolved the true ambiguities, and appropriate tuning values were determined for the Kalman filter. The system was declared ready for flight test.

4.4 Flight Tests

Seventeen flight test data runs were performed for a total of approximately 14 hours. The algorithm did not crash a single time during the 14 hours of flight testing. Even when messages were “lost”, satellites dropped out, the base satellite changed, or cycle slips occurred, the algorithm handled all the contingencies. There were a total of 91 satellite dropouts during actual flight testing. No dropouts while the aircraft were on the ground are included since these dropouts were due to missed messages (the antennas did not have clear line-of-sight). Of the 91 dropouts, 2 were due to erroneously detected cycle slips, 6 were actual cycle slips, and the remaining 83 dropouts were due to masking. Also, the base satellite changed 10 times over the 14 hours of flight test. Table 4.6 summarizes the satellite information. The algorithm

correctly handled every contingency, except when long strings of messages were “lost”, causing erroneous cycle slips on all of the satellites.

Table 4.6: Satellite Summary for Flight Testing

Total Data Sets	17 sets
Total Hours	14 hours
Base Satellite Changes	10 changes
Total Satellite Dropouts	91 dropouts
Erroneous Cycle Slips	2
Actual Cycle Slips	6
Satellites Masks	83

Not only did the algorithm properly handle contingencies, it also had good performance. The minimum indicator met the centimeter-level requirement in 13 of the 17 data runs, four of which were in the millimeter range. The times when the minimum indicator did not meet the centimeter-level requirement, it incorrectly chose an erroneous ambiguity set for a portion of the run. Table 4.7 gives a summary of the DRMS and MRSE for the MMAE and minimum indicator on each run. Runs in bold are times when the minimum indicator did not meet the centimeter requirement. Case F4.1 was the only run the that the MMAE performed better than the minimum indicator. The reason is discussed later in Section 4.4.4. Also, an average is shown in the last row.

Table 4.7 is a good summary of the errors during each data run. However, it does not tell the entire story. The DGPS system designed in this thesis was for use when the aircraft were in close formation (within 500 feet). Although interesting from an academic stand point, the times when the aircraft were outside this range should not be included in the error. For this reason, Figure 4.17 was constructed which displays every epoch when the aircraft were within 500 feet of each other. Each data run is separated by a black line with the appropriate date and time annotated on the plot. The x-axis shows the time from each data run (i.e., a time referenced within a particular data run in Figure 4.17 matches the time in the minimum indicator plots in

Table 4.7: Flight Test Position Error Summary (m)

Data Set	Min. Ind.		MMAE	
	DRMS	MRSE	DRMS	MRSE
F1.1	0.274	0.775	0.837	1.333
F1.2	0.041	0.085	0.363	0.417
F2.1	0.044	0.062	0.116	0.121
F2.2	0.004	0.023	0.343	0.738
F3.1	0.135	0.137	0.205	0.246
F3.2	0.090	0.153	0.192	0.229
F4.1	0.179	0.243	0.055	0.068
F5.1	0.002	0.011	0.004	0.013
F5.2	0.006	0.050	0.059	0.131
F6.1	0.003	0.010	0.074	0.138
F6.2	0.002	0.010	0.021	0.023
F6.3	0.004	0.007	0.012	0.021
F6.4	0.002	0.007	0.021	0.027
F7.1	0.002	0.008	0.011	0.015
F7.2	0.002	0.007	0.030	0.037
F8.1	0.003	0.012	0.054	0.072
F8.2	0.006	0.016	0.017	0.023
Average	0.047	0.095	0.142	0.215

Appendix A). The 1103L data run on 27 April was the only set in which the aircraft were not within 500 feet.

Over the course of approximately 11 hours, the DRMS and MRSE for the combined data runs was 1.9 and 3.3 centimeters, respectively! There was only one 40 second time period when the minimum indicator was not using the true ambiguity set. This was due to intermittent datalink problems which caused a system reset during the 1020L data run on 05 October data run. Excluding this 40 second period from the data, results in a remarkable DRMS of 0.3 centimeters and MRSE of 1.0 centimeter. The system performed almost flawlessly in its designed operational envelope.

The first flight data set, Case F1.1, is shown in its entirety, since it had the highest minimum indicator MRSE. Otherwise, only interesting data from specific runs

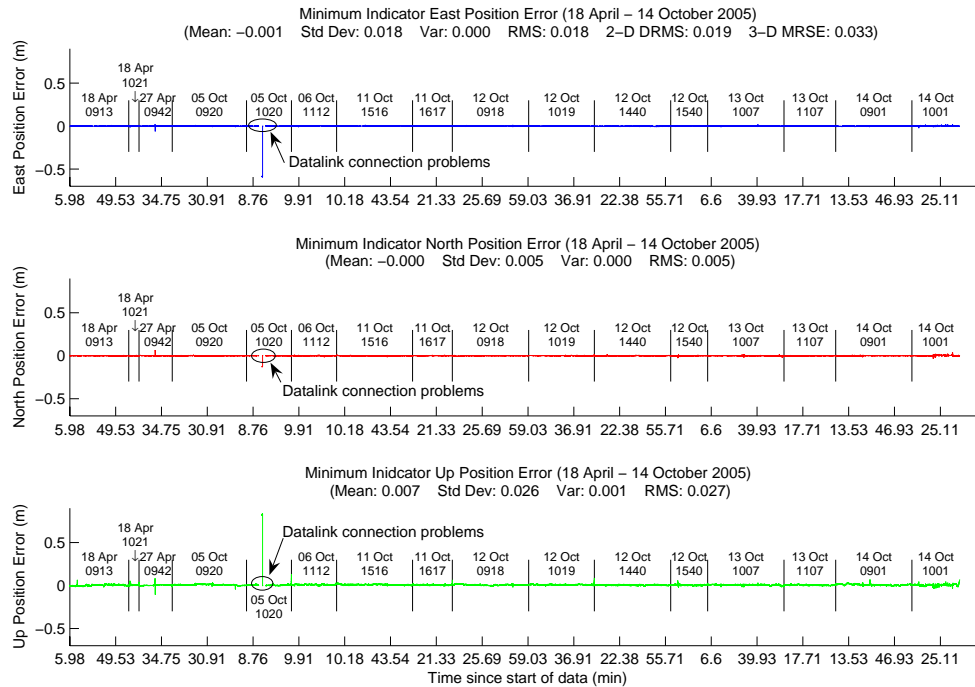


Figure 4.17: Every Epoch When the Aircraft Were Within 500 Feet

are displayed. The first four cases were from the April flight tests and the remaining sets were from the October flight tests. The first set of tests in April were between two C-12Cs. These tests verified the data link capability and DGPS algorithm position solution. The second set of tests in October was between a C-12C and Learjet 24. The Learjet had a modified flight control system that allowed the test team to input its own flight control laws for simulated aerial-refueling. The non-inertial hardware was identical between the April and October tests. Each flight test was divided into two data runs of approximately one hour each. Due to hardware problems, cases F2.1 and F2.2 were only 43 and 23 minutes long, respectively.

4.4.1 Floating Filter Reset. The data for this example came from Case F1.1, flown at EAFB on 18 April 2005. The flight took off at 0924L, and the weather was clear with some gusty winds. The actual data set begins just after engine start, runs through take-off and into the area work. During the test, the aircraft ranged from 18.8 to 1448.2 meters apart. Figure 4.18 shows the minimum indicator North, East, Down relative positions. The data run went well except for take-off.

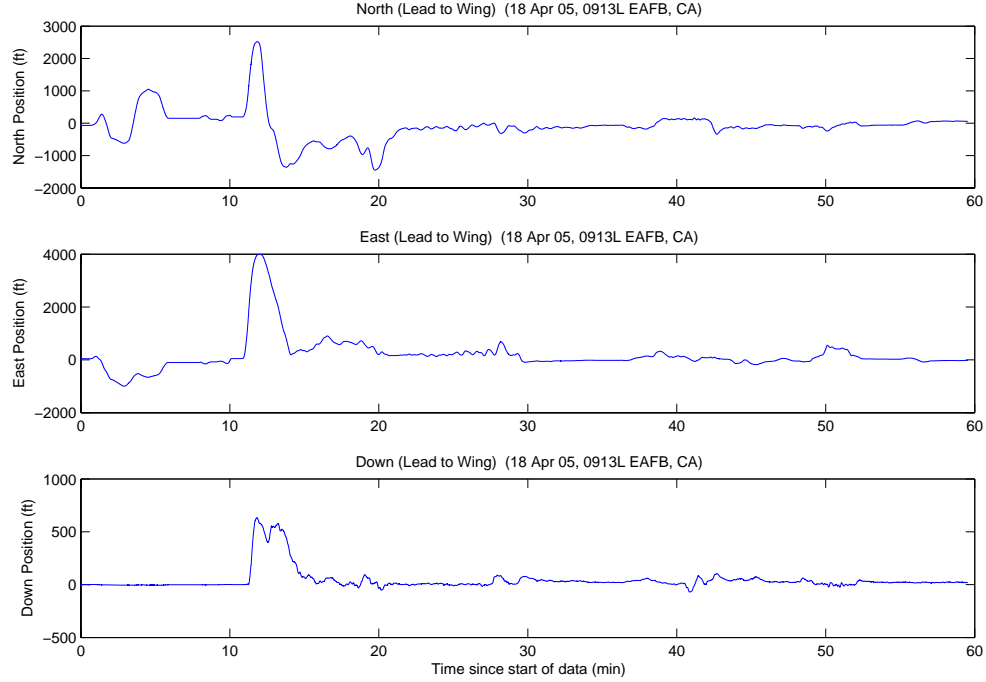


Figure 4.18: Case F1.1: North, East, Down Relative Positions

For an unknown reason, the floating filter reset approximately 50 feet above the ground during take-off. The lead aircraft had just begun a turn, masking the data link antenna from the wing aircraft. During post-processing, the floating filter did not reset, however five messages from the lead aircraft were lost at the exact time of the in-flight filter reset. Most likely, the algorithm incorrectly detected cycle slips on all of the satellites and set a reset flag.

It appears the missed messages caused the filter reset versus an algorithm problem. Also, the data link hardware was designed for line-of-sight operations only. For these reasons, the final summary table includes lines to show what the floating filter, minimum indicator and MMAE errors would be without the filter reset.

The LAMBDA function returned a total of 29 unique ambiguity sets during the data run. The correct set was the top candidate returned from the first LAMBDA pair. The minimum indicator correctly chose the true ambiguity set at four minutes. It did not change again until after the floating filter reset. The true ambiguity set was rediscovered at 15.16 minutes, after two other incorrect ambiguity sets were chosen.

After the filter reset, the true ambiguity set was returned from LAMBDA as the number one set, 2732 of 2876 epochs (95.0 percent), with a maximum ratio of 33.1.

Given the filter reset, the errors in the floating filter, minimum indicator and MMAE were larger than other data runs without the filter reset. If a filter reset occurred, the entire system was reset to include the floating filter, MMAE, and minimum indicator. Basically, the system was re-initialized. Figure 4.19 shows the floating filter position errors.

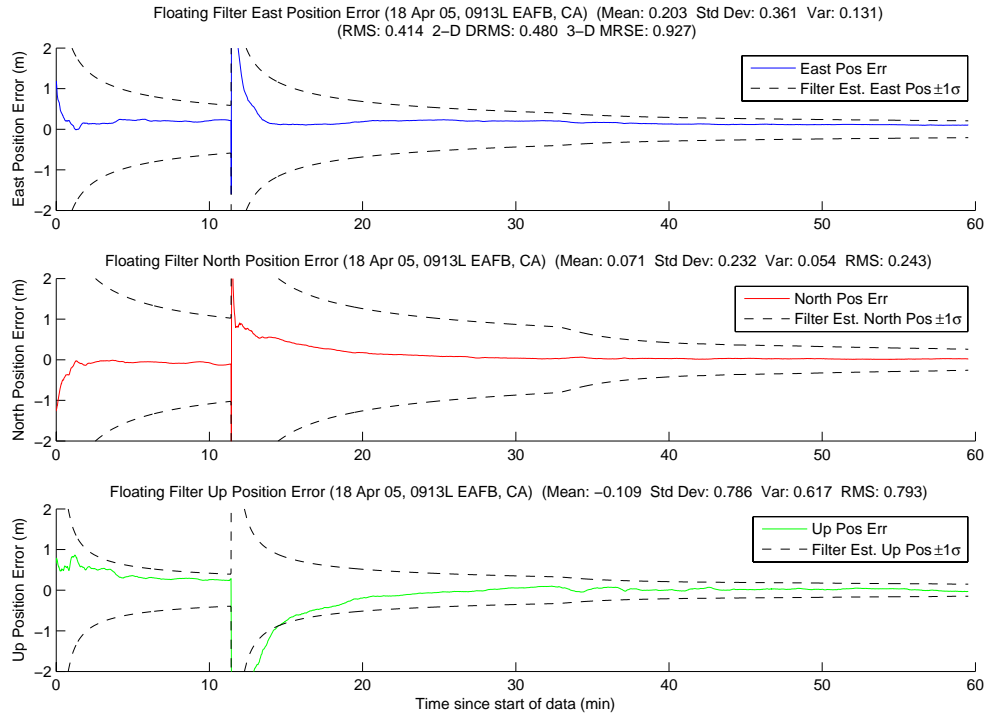


Figure 4.19: Case F1.1: Floating Filter Position Errors

Figure 4.20 displays PRN 6 measurement information and the floating filter ambiguity estimate. The extremely large phase residuals at approximately 11 minutes coincide with the filter reset. The static initialization and other times when both aircraft were stopped can be seen in the phase residual plot at zero, seven and ten minutes.

Figure 4.21 shows the number of visible satellites. Most satellite dropouts were due to aircraft maneuvering. The C-12C had high mounted horizontal stabilizers and

the GPS antenna was mounted on the center of the tail. When the aircraft banked, the horizontal stabilizer could mask certain satellites depending on the satellite geometry and bank angle. Since the elevation cutoff angle for satellites was 10 degrees, bank angles as small as 10 degrees could potentially mask satellites. It was not uncommon for the aircraft to bank 30 degrees or more during a data run. Even though the filter initialized with six satellites, the aircraft were static on the ramp with engines running, and the floating filter was able to converge normally.

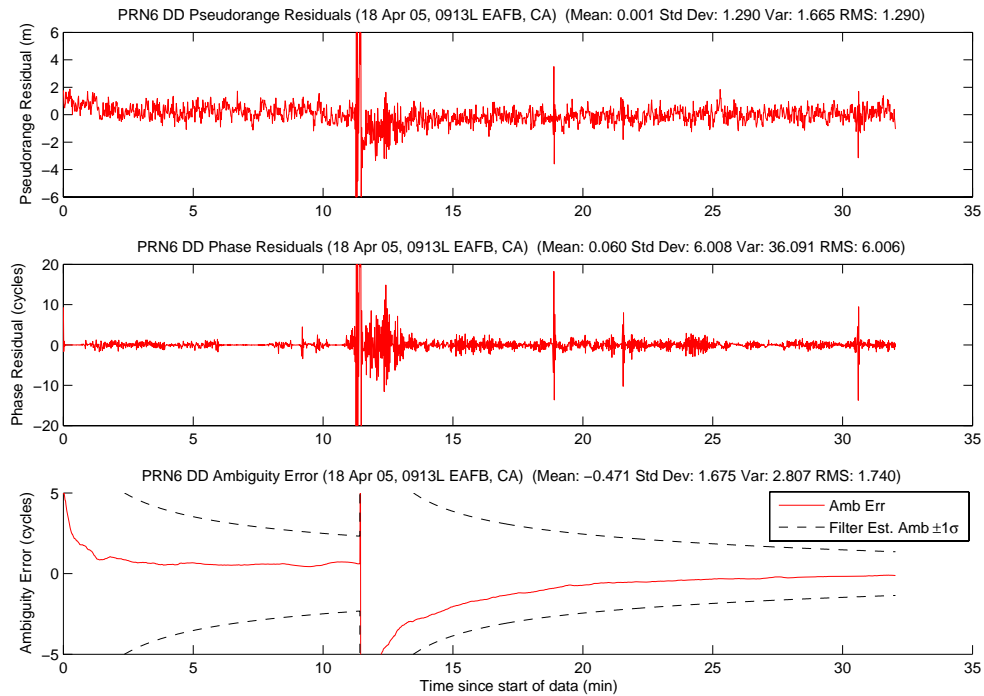


Figure 4.20: Case F1.1: Satellite 6 Measurements

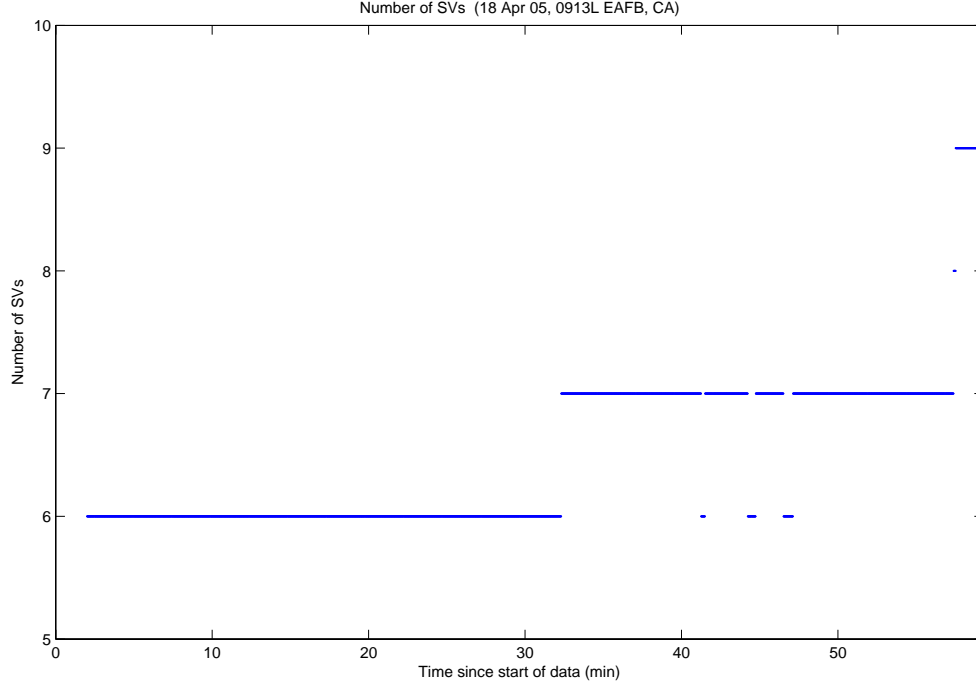


Figure 4.21: Case F1.1: Visible Satellites

Once more, the residual plot, Figure 4.22, is expanded to help identify the true ambiguity set's residuals. The filter reset at 11 minutes is clearly visible with a “wall” of erroneous ambiguity sets' residuals. The probability plot, Figure 4.23, has one point of interest. The filter reset occurred at 11 minutes when all the probabilities simultaneously drop to zero. Interestingly, the original EF[1] and what eventually became EF[15] have matching ambiguity sets. As the LAMBDA algorithm repopulated the ambiguity bank, an incorrect set initially absorbed 50 percent of the probability. Once the true set was returned from the LAMBDA function, it absorbed all of the probability by 26 minutes.

The minimum indicator discovered the true set, the second time, at 15.8 minutes (after selecting two wrong sets). The precise time the true ambiguity set was returned from the LAMBDA algorithm was not known. Assuming the minimum indicator immediately selected the true set after the initial two-minute SOS residuals acquisition delay, the true ambiguity set was returned from LAMBDA at 13.8 minutes (approx-

imately three minutes after the reset). After the minimum indicator reselected the true set, it did not change during the remainder of the data run.

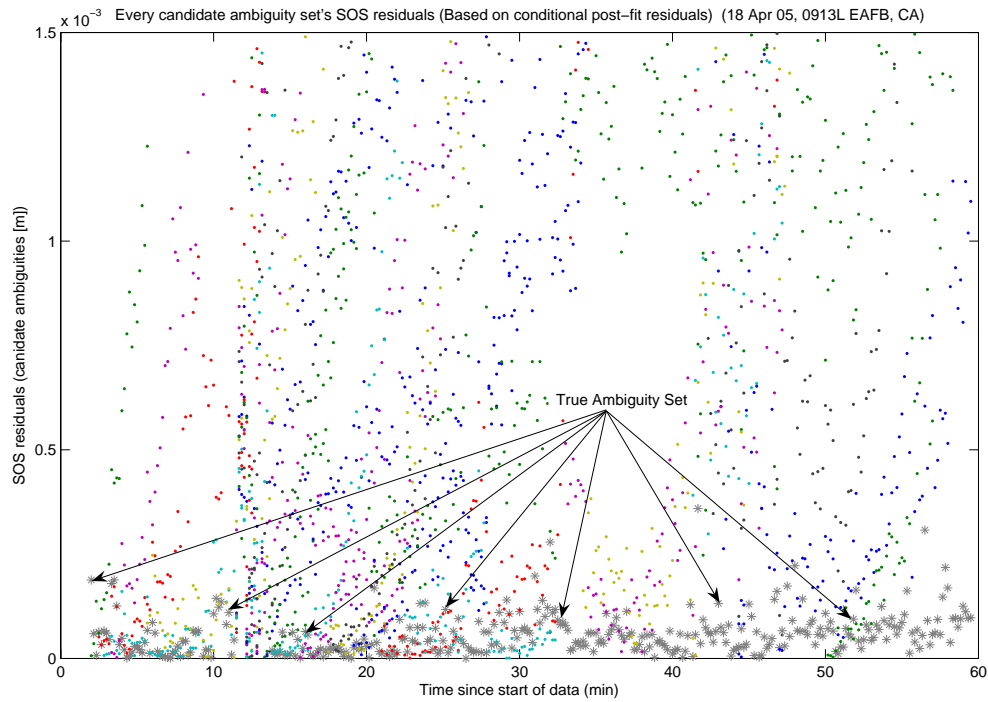


Figure 4.22: Case F1.1: SOS Residuals

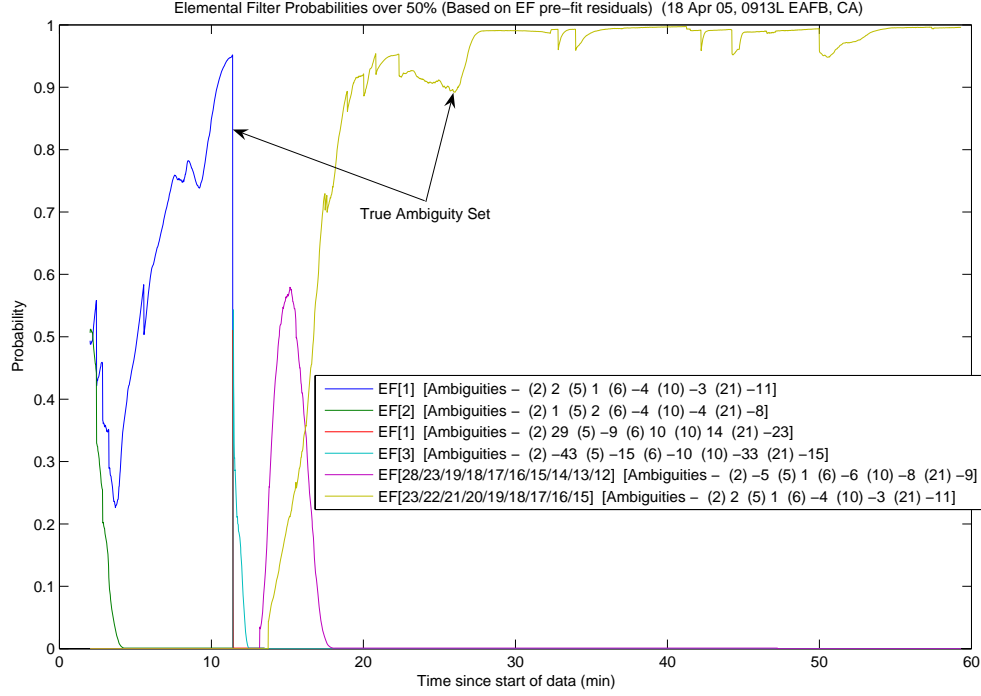


Figure 4.23: Case F1.1: Conditional Probabilities for Selected MMAE Elemental Filters

The MMAE position error is shown in Figure 4.24. The majority of the error was from the filter reset at 11 minutes. Smaller divergences in the MMAE position error due to probability fluctuations are not visible.

Finally, Figure 4.25 displays the position error plot for the minimum indicator variable, and Table 4.8 gives a summary of each method's position error. Table 4.8 also includes the post-processing error numbers since the filter did not reset while in post-processing. Just like the MMAE position errors, the minimum indicator error also suffered from the filter reset. The error immediately dropped at 15.8 minutes when the true ambiguity set was reacquired. There were no minimum indicator values from the floating filter reset at 11 minutes, until the first incorrect ambiguity set was selected at approximately 13 minutes. Therefore, there is a gap in the minimum indicator during those times.

Unfortunately, the minimum indicator error was not within the desired centimeter-level accuracy for MRSE with the floating filter reset. Of course, without the

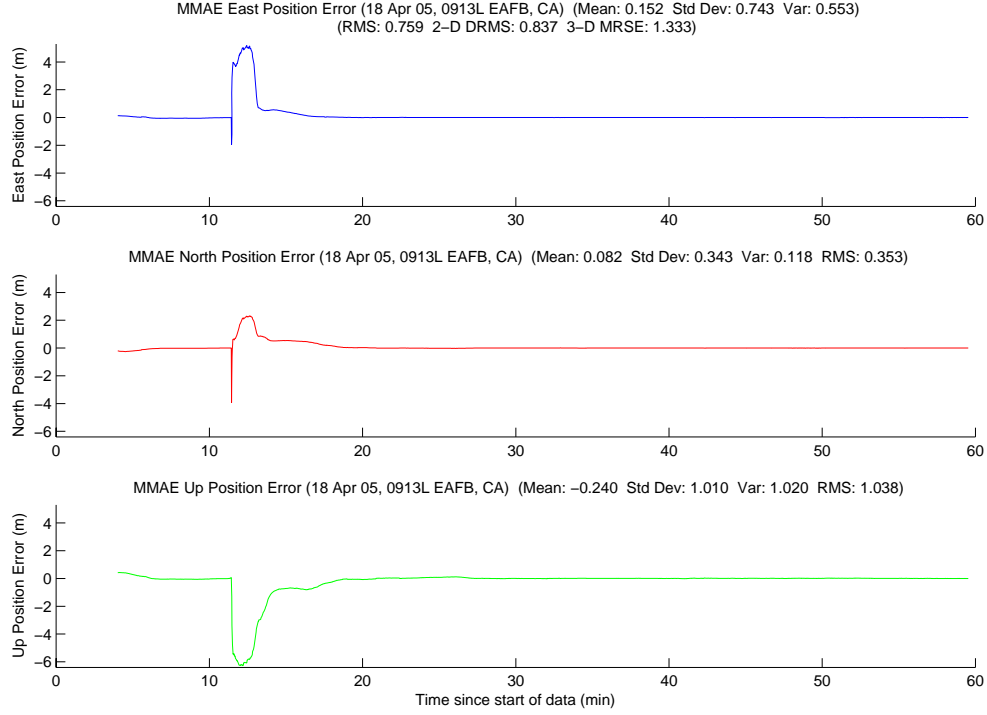


Figure 4.24: Case F1.1: MMAE Position Error

filter reset, the minimum indicator remained well within the tolerance and was similar in magnitude to the dynamic ground test results. Most importantly, the algorithm and hardware performed extremely well during the flight for which the system was designed to run (i.e., close formation). In every flight test run, when the aircraft were in close formation, the algorithm was within the centimeter-level requirement.

Table 4.8: Case F1.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.203	0.361	0.071	0.232	-0.109	0.786	0.480	0.927
MMAE	0.152	0.743	0.082	0.343	-0.240	1.010	0.837	1.333
Min. Ind.	0.030	0.170	0.039	0.208	-0.116	0.716	0.274	0.775
Float Filter*	0.165	0.060	-0.046	0.087	0.169	0.144	0.201	0.299
MMAE*	0.001	0.024	-0.010	0.037	0.020	0.065	0.045	0.082
Min. Ind.*	-0.000	0.001	-0.000	0.001	0.002	0.008	0.002	0.009

(Note: * are the values without the floating filter reset at 11 minutes.)

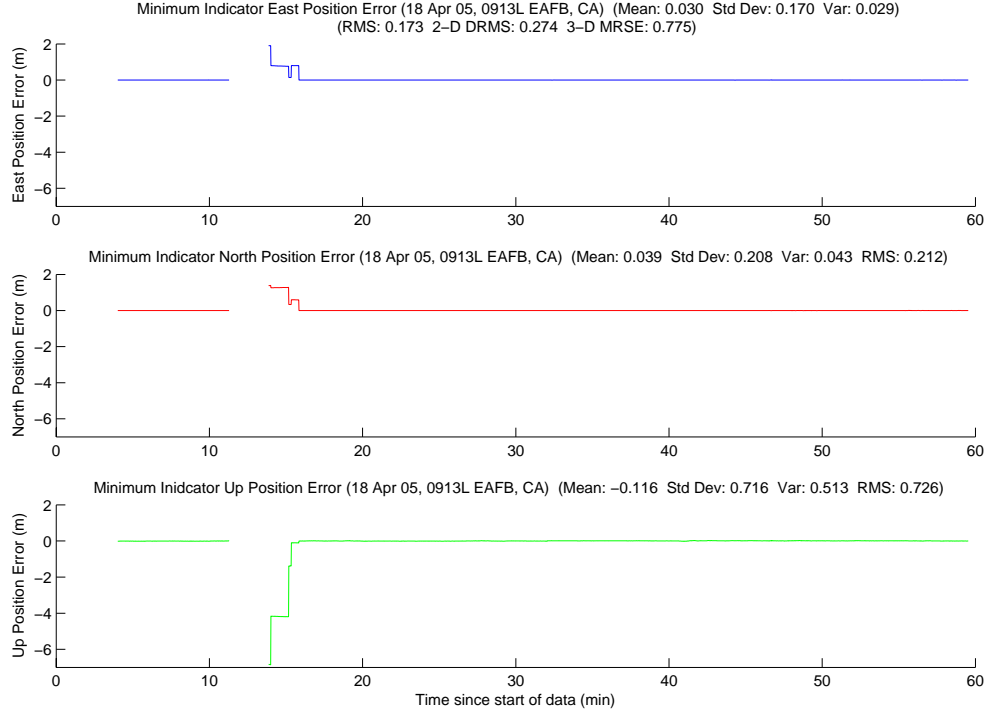


Figure 4.25: Case F1.1: Minimum Indicator Position Error

Case F1.1 is the last case shown in entirety. The following examples show interesting events from selected cases during flight testing. The examples are still broken down by case numbers. The first example is from Case F1.2; it shows what happens when the true ambiguity set is not in the candidate ambiguity bank.

4.4.2 True Set Not in Candidate Ambiguity Bank. The data for this example came from Case F1.2. It started eight minutes after the end of Case F1.1. Some very aggressive maneuvering by both aircraft at 52 minutes into the data run caused multiple cycle slips. Both aircraft were banking over 45 degrees in either direction and pulling approximately 2.5 gs. Due to the cycle slips, the minimum indicator was wrong during the last eight minutes of the data run. There was not enough time for the LAMBDA function to return the true set before the end of the data run. More discussion about the true ambiguity set follows with the SOS residuals plot.

The SOS residuals plot, Figure 4.26, shows some interesting results. The true ambiguity set's SOS residuals were easily detectable up until the cycle slips at 52

minutes. After the cycle slips, no set appeared to “fill the gap” left by the true ambiguity set. This was the tale-tell sign that the true ambiguity set was not in the ambiguity bank. The probability plot, Figure 4.27, reveals the same information. Numerous ambiguity sets absorbed over 99 percent of the conditional probability during the last eight minutes of the run. Since the true ambiguity set was not in the ambiguity bank, neither the minimum indicator nor the MMAE could correctly identify the true set. However, even with the cycle slips, the minimum indicator MRSE achieved the centimeter-level requirement. The next case is a good example of a local minimum for an erroneous ambiguity set.

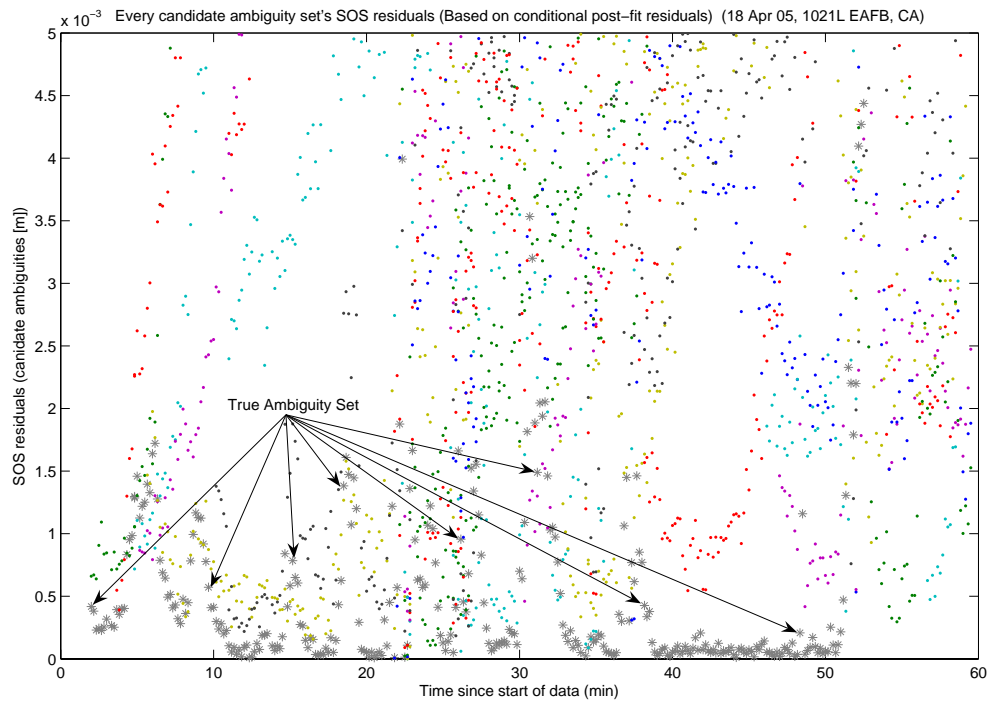


Figure 4.26: Case F1.2: SOS Residuals

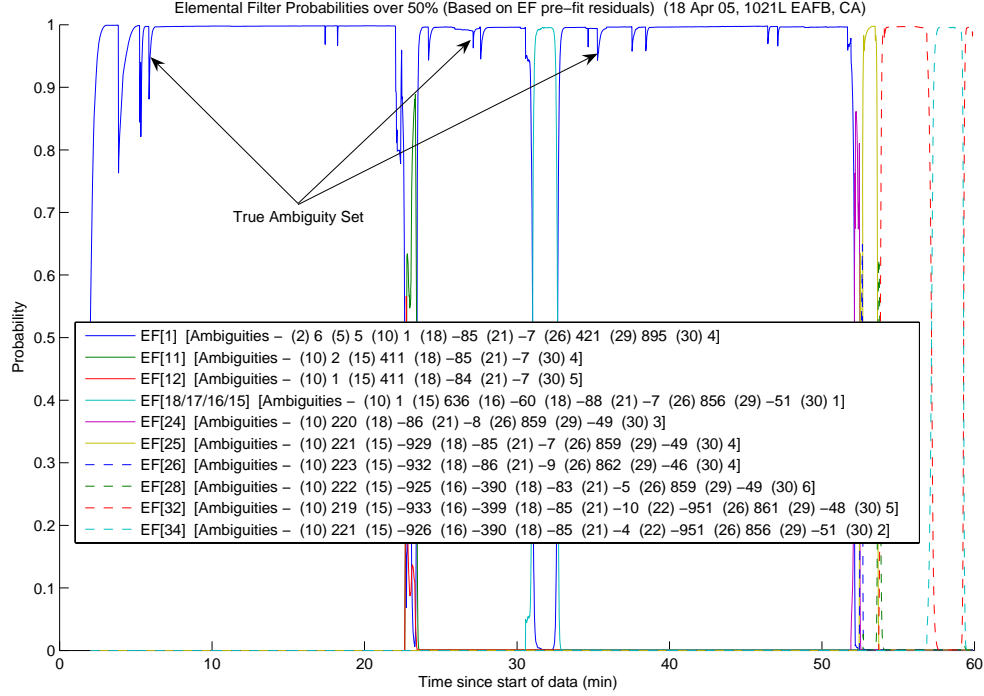


Figure 4.27: Case F1.2: Conditional Probabilities for Selected MMAE Elemental Filters

4.4.3 Erroneous Ambiguity Set Local Minimum. The data for this example came from Case F2.2. This was the second data set of the second test flight, flown on 27 April 2005. The data run was terminated early due to hardware problems with the Javad® receiver in the trail aircraft. Therefore, some of the information was taken from post-processed data. The SOS residuals plot and conditional probability plots showed an unusual occurrence.

The SOS residuals plot and probability plot were constructed from post-processed data. Figure 4.28 shows all of the candidate ambiguity sets' SOS residuals. The increase in the true set's SOS residuals at 11.5 minutes occurred when the aircraft were maneuvering 1.4 km apart. Figure 4.29 displays the MMAE conditional probability. The true ambiguity set maintained over 49 percent of the total probability until 13 minutes. Then, an erroneous set slowly absorbed 99 percent of the probability. Once the wrong ambiguity set reached 99 percent at 15 minutes, the true set slowly began reabsorbing the probability back. An examination of the SOS residual plot shows

why this phenomenon occurred. The erroneous set's SOS residuals are highlighted by gold open squares in Figure 4.28. The erroneous set's residuals remain lower, on a parallel course, to the true set's residuals starting at approximately 13 minutes. Once the wrong set's residuals start growing again at 15 minutes, and the true set levels off, the true ambiguity set begins to reabsorb the probability. This was a classic example of how an incorrect ambiguity set can have a local minimum below the true ambiguity set. The next example, Case 4.1, is discussed since it was the only test case in which the average MMAE position solution was better than the average minimum indicator position solution.

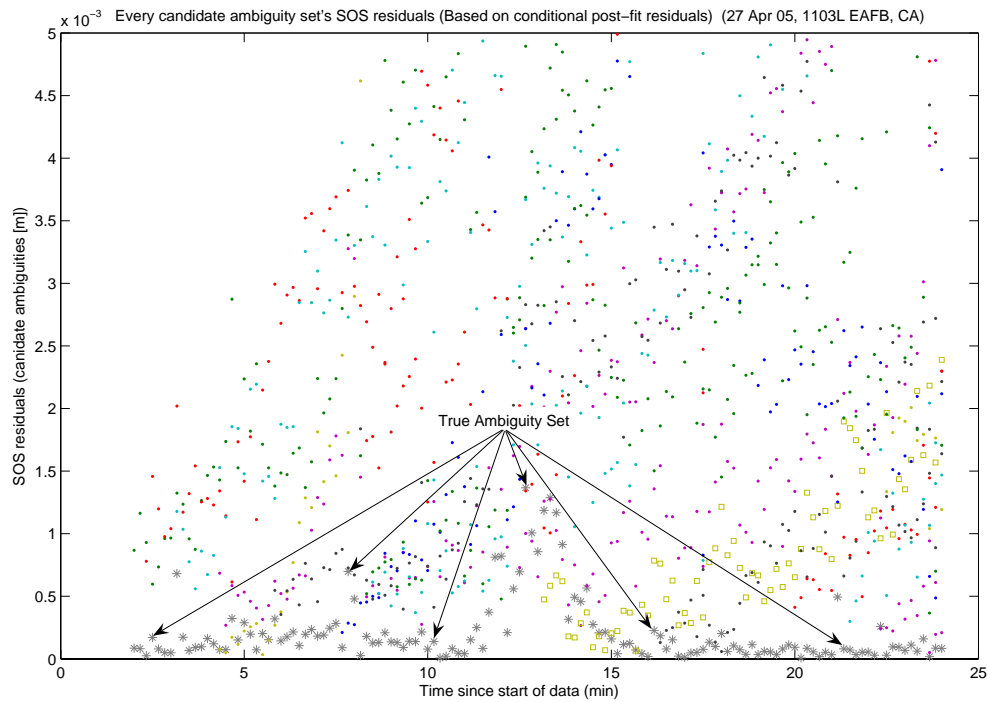


Figure 4.28: Case F2.2: SOS Residuals

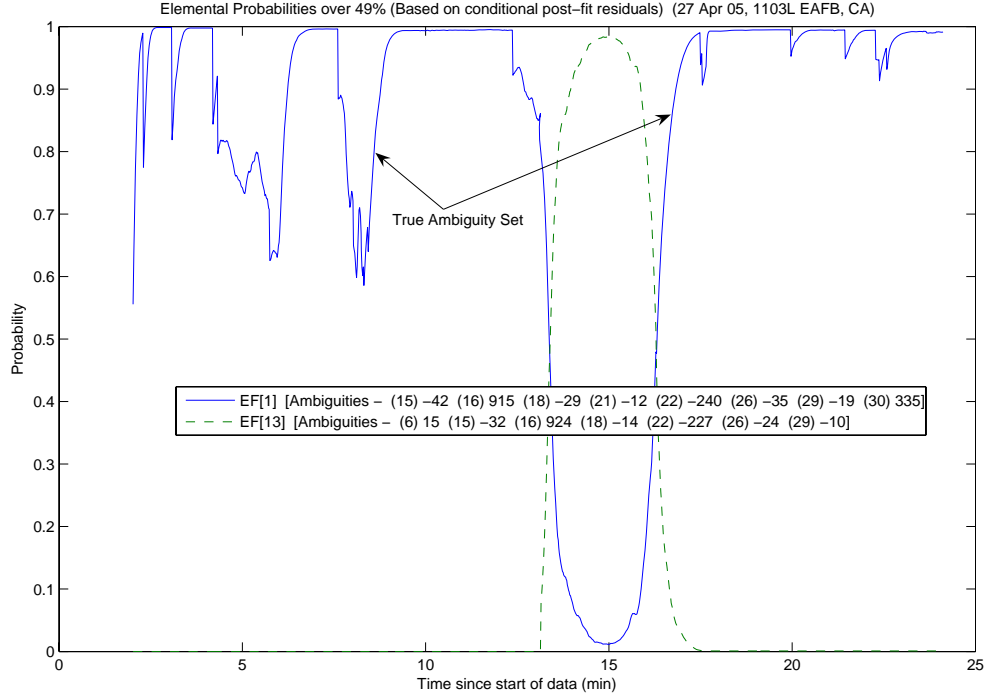


Figure 4.29: Case F2.2: Conditional Probabilities for Selected MMAE Elemental Filters

4.4.4 MMAE Outperforms Minimum Indicator. The data for this example came from Case F4.1, flown on 6 October 2005. The flight terminated early due to problems unrelated to the test hardware or software described in this thesis. Therefore, there was only one data run collected and it was 38 minutes long. Case F4.1 was the only data set in which the average MMAE position solution was better than the average minimum indicator position solution.

The system was started once airborne when the aircraft were 1530 meters apart. The true ambiguity set was the third set returned from the LAMBDA function, 47 seconds after the first ambiguity set was returned. The minimum indicator initially chose an incorrect ambiguity set at four minutes. It changed to the true ambiguity set 50 seconds later. During the 50 second period when the minimum indicator chose the wrong set, the total error was approximately two meters, as seen in Figure 4.30. It was this 50 second period of two meter error that drove the minimum indicator MRSE above the MMAE MRSE. However, the aircraft were 1500 meters apart during

this time. Once the aircraft joined, for the critical part of the flight, the error was well within the centimeter requirement. It is also interesting to note that, with the true ambiguity set, a position accuracy under 0.5 meters was possible when the aircraft were 3.5 kilometers apart. This is evident by the small “tail” at the end of the North and Up subplots when the aircraft were separating for return to base (RTB).

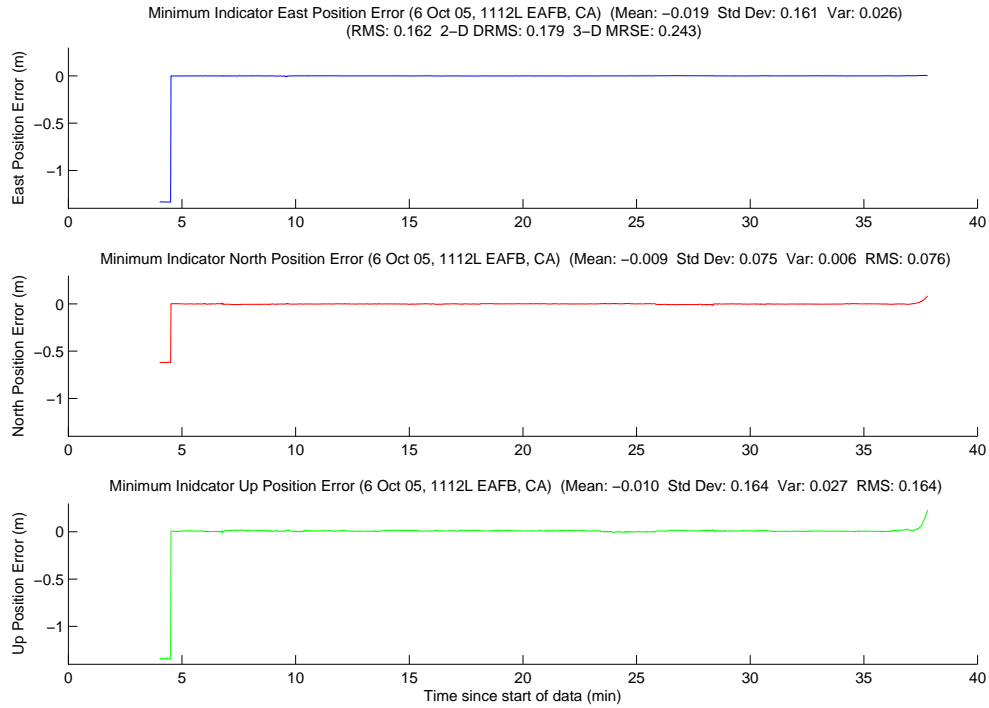


Figure 4.30: Case F4.1: Minimum Indicator Position Error

In contrast, the MMAE position error plot in Figure 4.31 is relatively accurate, even at the start of the data run. The true ambiguity set had absorbed approximately 90 percent of the probability by four minutes into the run. Since the MMAE position solution is weighted, and the true ambiguity had absorbed almost all of the probability by four minutes, the average MMAE position solution was more accurate. However, between 5 and 11 minutes, and at the end of the data run, the minimum indicator was an order of magnitude more accurate than the MMAE. The next example will show the power of the number of visible satellites.

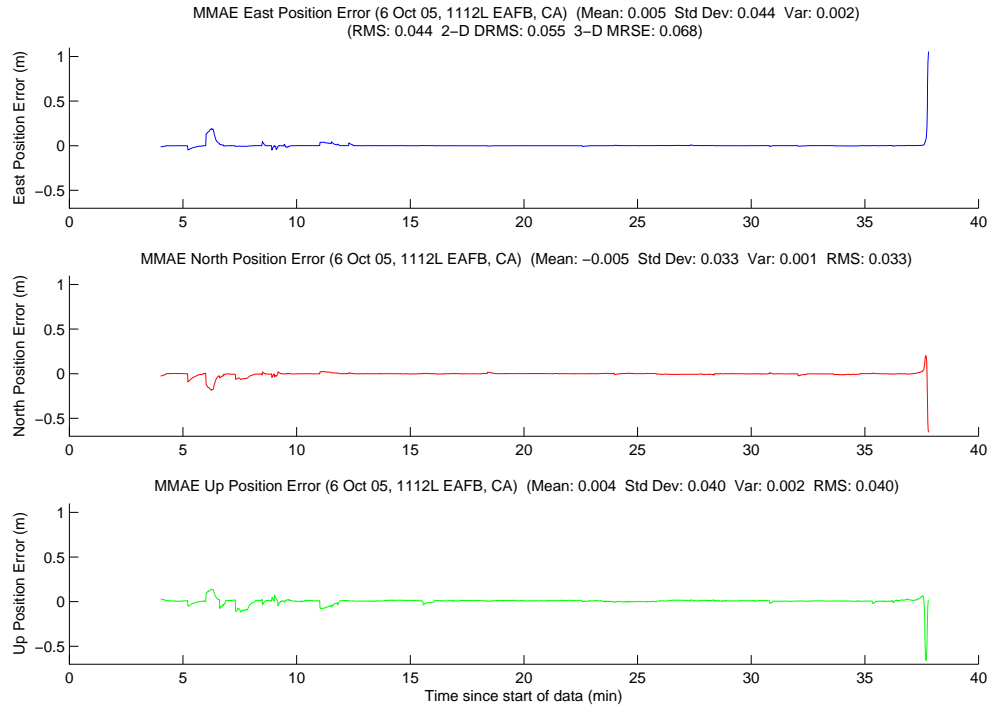


Figure 4.31: Case F4.1: MMAE Position Error

4.4.5 Twelve Visible Satellites. The data for this example came from Case F5.1, flown on 11 October 2005 at 1516L. Case F5.1 is interesting because of the number of visible satellites. The data run started with 10 satellites. Then, three minutes into the run, another satellite rose into view for 11 total. For approximately eight minutes in the middle of the data run, there were 12 visible satellites. The Javad® JNS100 receivers only had 12 channels, so that was the maximum number of useable satellites. Figure 4.32 shows the number of visible satellites. The reason the number of satellites is so interesting is apparent in the SOS residual and conditional probability plot.

Figure 4.33 shows all of the candidate ambiguity sets' SOS residuals. The importance of the number of satellites is dramatically evident in the residual plot. From the beginning of the data run, there was not a single erroneous set's residuals near the true set. Even without the emphasis, the true set's residuals are clearly visible. As the number of satellites increased, it became much more difficult for an incorrect ambiguity set's measurements to "fit" the position solution. Starting at

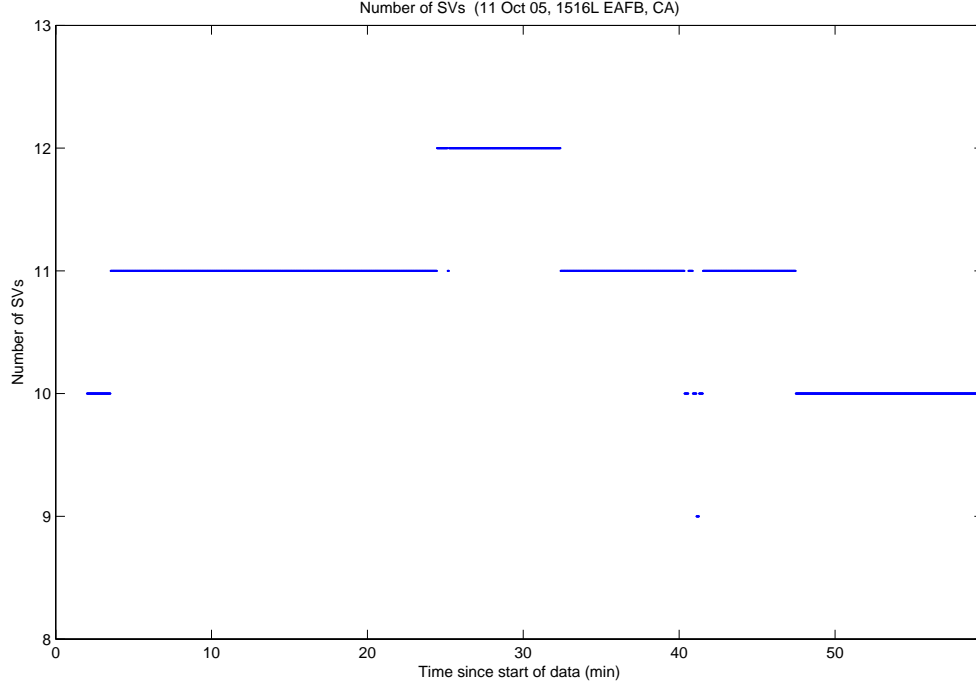


Figure 4.32: Case F5.1: Visible Satellites

approximately 25 minutes into the data run, when there were 12 visible satellites, no incorrect set's residuals are visible in the plot until approximately 40 minutes. Figure 4.33 is similar to the widelane SOS residuals plot shown later in the chapter. Figure 4.34 displays the MMAE conditional probability. The true ambiguity set maintained over 49 percent of the total probability for the entire data run. After the first two dips, the true set's probability did not drop below 90 percent.

The final example comes from the last flight data set, Case F8.2. Case F8.2 shows the robustness of the algorithm in the face of numerous satellite dropouts.

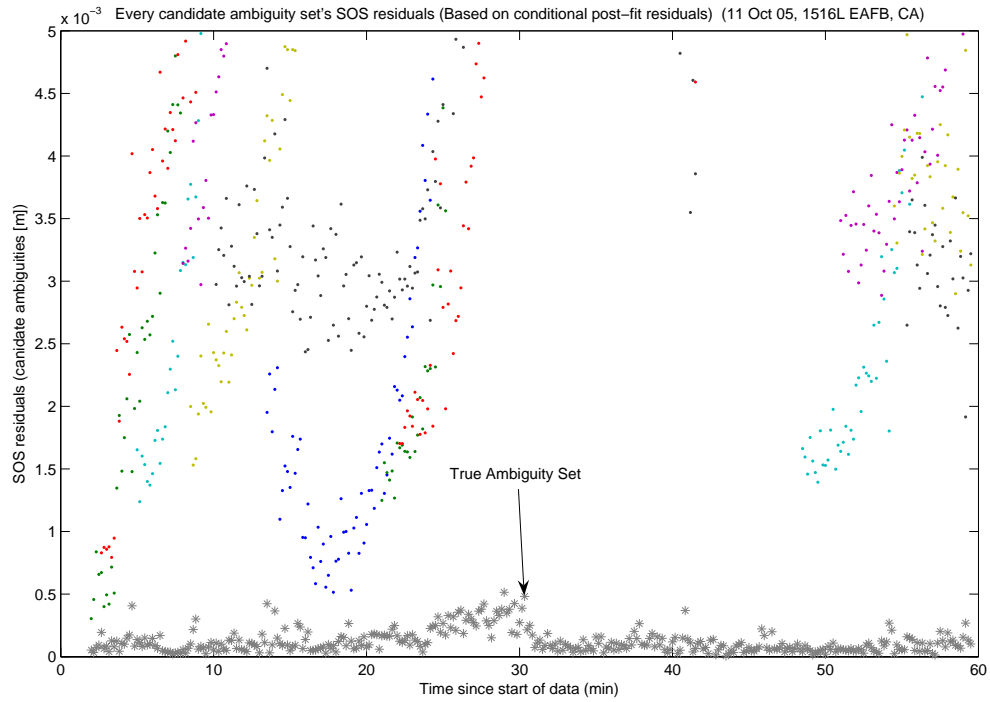


Figure 4.33: Case F5.1: SOS Residuals

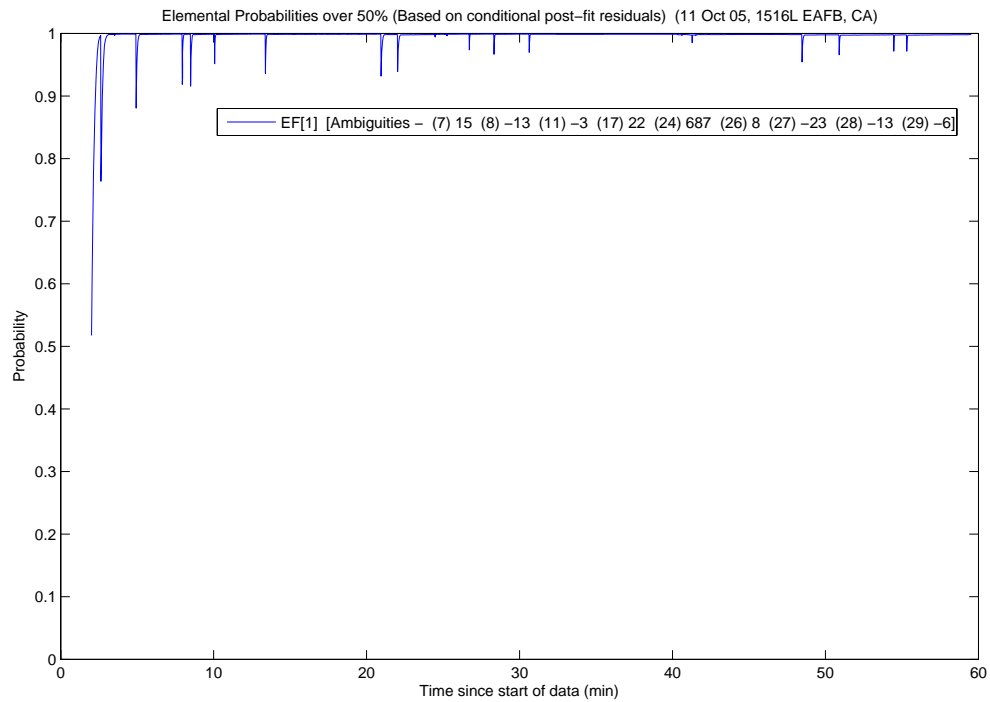


Figure 4.34: Case F5.1: Conditional Probabilities for Selected MMAE Elemental Filters

4.4.6 Multiple Satellite Dropouts. The data for this example came from Case F8.2, and it was the final flight test data run. It started at 1001L on 14 October 2005. Figure 4.35 shows the number of visible satellites during the run. There was a substantial amount of satellite activity during the data run. There were 28 dropouts, with no satellites rising or setting. Of the 28 dropouts, three were due to cycle slips. Each cycle slip was correctly identified by the triple-difference method discussed in Section 3.9.1. The cycle slips occurred on PRNs 3, 8, and 20 at 9, 14, and 26 minutes, respectively. PRN 8 had the most dropouts with 10. The rest were split fairly evenly between PRNs 16, 19, and 20. The 28 dropouts were the most dropouts of any data run. The most likely cause for the increase in dropouts and cycle slips was a third aircraft in the formation taking video. The floating filter, MMAE, and minimum indicator still performed very well, a testament to the algorithm’s robustness.

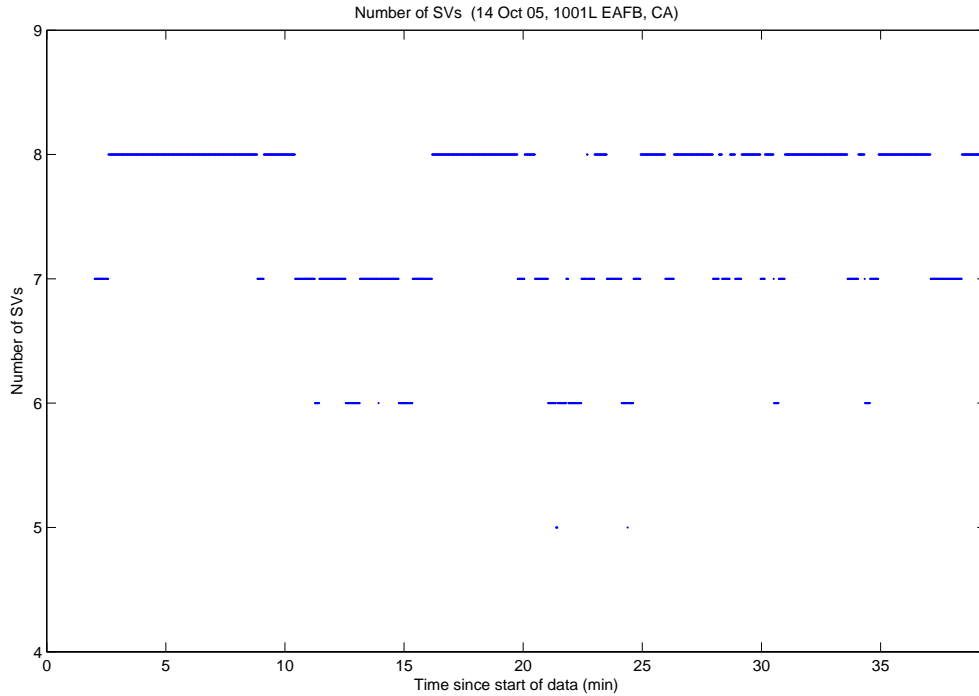


Figure 4.35: Case F8.2: Visible Satellites

Finally, Figure 4.36 displays the position error plot for the minimum indicator. The position errors on the minimum indicator had numerous small jittery spikes, especially from 10 to 15 minutes, and from 20 to 35 minutes. These times corre-

sponded to when there were numerous dropouts and cycle slips. Phase measurements from the Javad® receivers were noisy (possibly due to aircraft masking of satellites and potential multipath). The minimum indicator MRSE was still well within the centimeter-level requirement at 0.016 meters. Case F8.2 concludes the flight test results.

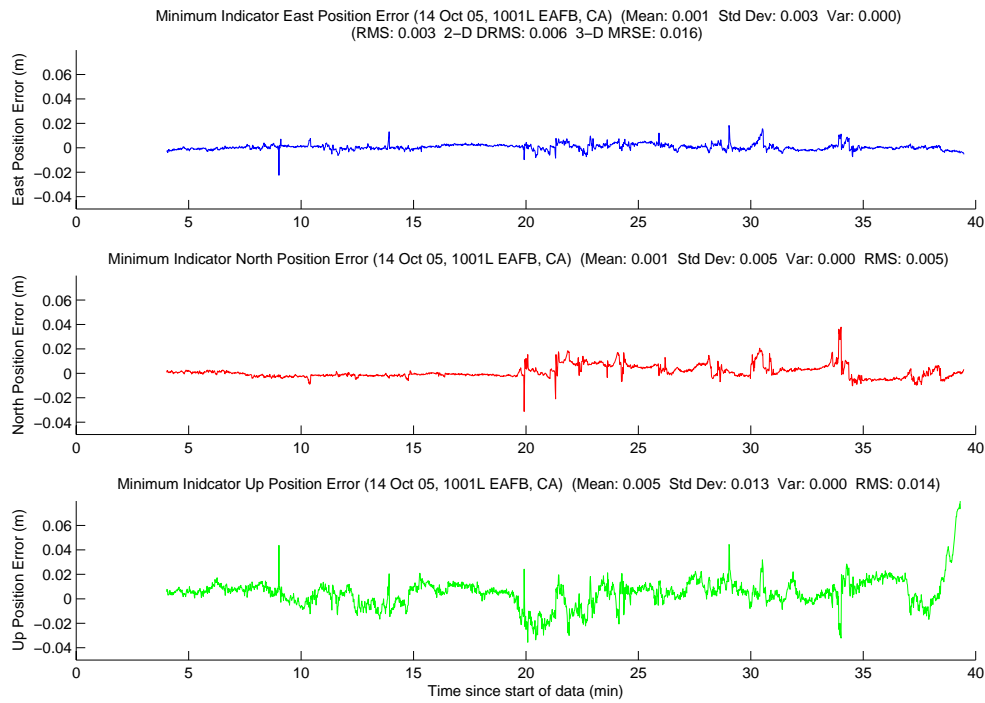


Figure 4.36: Case F8.2: Minimum Indicator Position Error

4.4.7 Real-time Timing. The algorithm developed in this thesis was designed to run in real-time. The floating filter was updated every second, and the precise relative position was output at a 20 Hz rate as discussed in Section 3.4. Two important aspects needed to be addressed for the system to operate correctly. First, the total delay from when a GPS measurement arrived until the precise relative position message was output to aircraft bus needed to be small as possible. On average, this delay was approximately 80 milliseconds, well within the desired tolerance.

The second aspect was the inter-message delay (i.e., the time between consecutive messages). The aircraft interface had strict timing guidelines. If the inter-message delay became too large, messages would be skipped. Prior to the October flight tests, no investigation was made on the inter-message delay. During the October testing, a one second anomaly in the precise relative GPS position data stream manifested itself in the loss of a single relative position message. It did not always occur, and there was seemingly no correlation to other factors when it happened. At first, the floating filter update was suspected based on its one Hz update rate. However, since the filter was running in a separate thread from the main program, this was a slim possibility.

During post-processing, it was discovered that the precise relative position messages were being sent at non-uniform intervals. Normally, the delay varied between 31 and 62 milliseconds. Sometimes the delay was 93 milliseconds. It was strange that the delay increased in 31 millisecond intervals. Further, after summing the inter-message delay for 20 messages, the total delay occasionally equaled 1.031 seconds. The total delay should be one second, and the additional 0.031 seconds was exactly one “normal” message delay period.

After some research, it was determined that the cause was most likely due to context switching on the Intel® processor between the multiple program threads. There were normally five threads running on the single processor in the trail aircraft. The first thread was listening for incoming messages from the lead aircraft. The second thread was the main program running. Third was the floating filter thread.

The fourth was the thread listing for the precise relative position message. Finally, there was a thread for message parsing / cyclic redundancy checks (CRC) and delivery routing. Generally, there is a 5-10 millisecond delay between threads on the processor used in this thesis. A few context switches could put the precise relative position message thread behind. A faster processor or multiple processors could help the situation.

4.4.8 Flight Test Summary. Six flight test examples were presented showing the effects of a floating filter reset and cycle slips, an erroneous ambiguity set local minimum, the power of the number of visible satellites, and the effect of multiple satellite dropouts. Additional plots from the flight tests are in Appendix A. The algorithm properly handled every unusual event. Unfortunately, there were some timing problems, but the system as a proof-of-concept worked well. Most importantly, all of the flight objectives were met.

The first objective was to determine the relative position accuracy. The average minimum indicator position error was typically within the centimeter-level requirement during the hour-long data runs. However, the minimum indicator was always within the centimeter-level requirement during all close formation flying (the critical part of the flights).

The second objective was to determine data link robustness. The data link operated reliably from fingertip formation (one meter wingtip clearance) out to approximately three kilometers. There was only one flight test in which the data link had problems in relatively close formation (30 meters) for an unknown reason.

The final flight objective was to determine IMU accuracy. This was accomplished by the “Lost Wingman” test team during the April 2005 flight tests [16]. The final section of Chapter IV discusses the special cases mentioned earlier.

4.5 Special Cases

Four special cases are presented. The first special case is comparing “pre-fit” and “post-fit” residuals. Next, the filter in position-only mode is shown, followed by the widelane measurement observables example. Finally, an example of the floating filter with the dynamic T_a and σ_a^2 option turned on is presented.

4.5.1 “Pre-Fit” and “Post-Fit” Residual Comparison. As mentioned previously in Section 3.6.2, it was originally thought that “post-fit” residuals were less “noisy” and would produce better results from the MMAE probability calculation. Figure 4.37 shows an example of the noisy “pre-fit” residuals, contrasted with the same “post-fit” residuals in Figure 4.38.

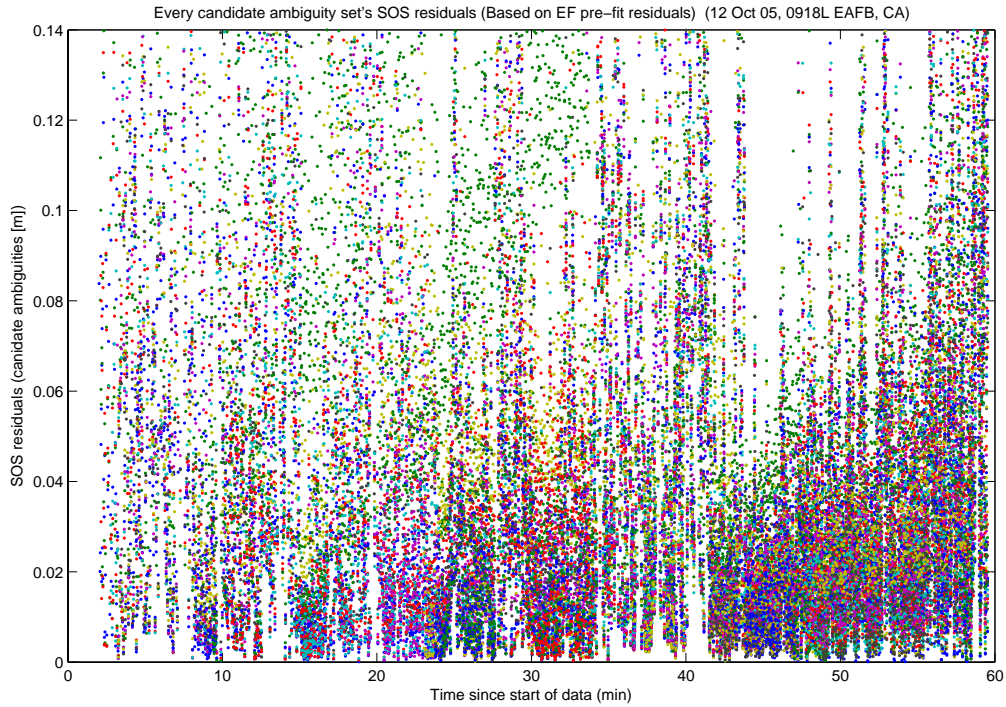


Figure 4.37: “Pre-fit” Residuals

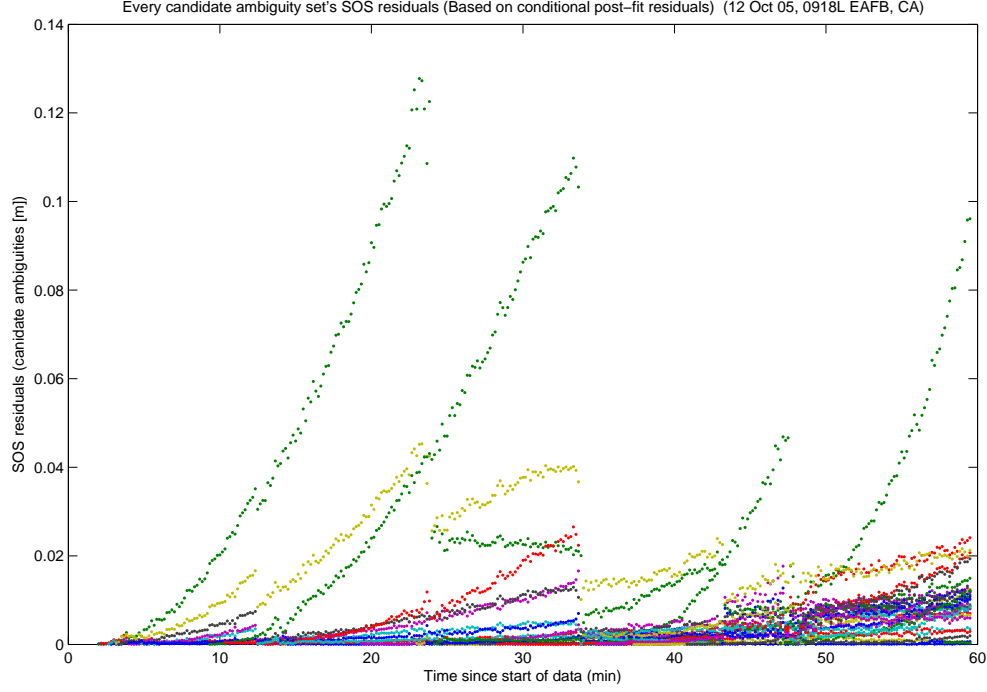


Figure 4.38: “Post-fit” Residuals

Even though it appeared the “post-fit” residuals would produce better results, Ormsby proved mathematically that the MMAE probability calculation was equivalent with either “pre-fit” or “post-fit” residuals [44]. To determine if this was true, the probability calculation was done with both “pre-fit” and “post-fit” residuals, and then the results were compared. Before the results are displayed, it is important to note that the comparison shown here is not exactly “apples-to-apples”. The “pre-fit” residuals came from the MMAE, but the “post-fit” residuals were from the epoch-by-epoch calculation described in Section 3.7.1. The real difference is that the MMAE “remembers” previous information garnered during a data run, whereas the epoch-by-epoch calculation has no “knowledge” of previous measurements. Therefore, the probability calculations will not be exactly the same, but should be similar. Only one comparison is shown here, but all of the comparison plots are in Appendix A.

Figure 4.39 is from the first data set of the first flight on 12 October 2005. The top subplot is the raw probability data overlaid on each other. The bottom subplot shows the difference between the two probabilities. As Ormsby proved, the

plot proves that the probability calculation is almost identical between the “pre” and “post-residuals”. The magnitude of the dip at 35 minutes was different by about 15 percent, however, the trend of each line is identical. Other probability plots from different test cases were similar to Figure 4.39. The position-only filter is discussed next.

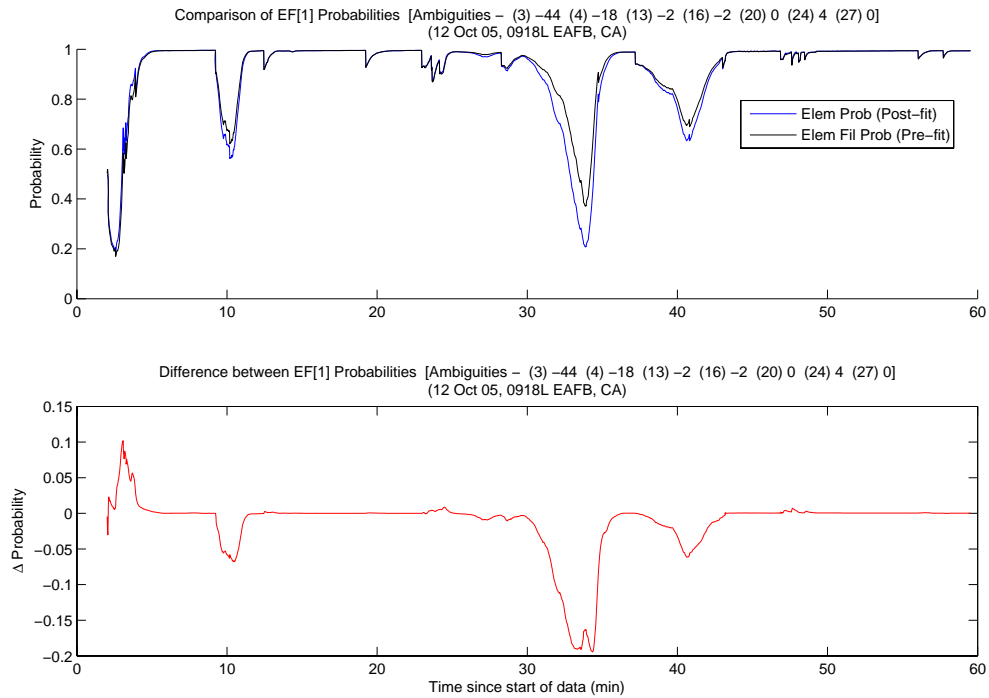


Figure 4.39: “Pre-fit”/“Post-fit” Residual Case: Probability Comparison

4.5.2 Position-Only Filter Comparison. Since the Kalman filter designed in this thesis was so heavily driven by the measurements and not the FOGMA model, the prospect of reducing computational complexity by removing the relative velocity and relative acceleration states was investigated. Also, since the filter was designed to run in real-time, a reduction in computational complexity translated directly to less time for floating filter and MMAE calculations. As a reminder, the relative position states, in the position-only filter, had a noise value of $200 \text{ m}^2/\text{sec}$ in the process noise matrix, \mathbf{Q} , as specified in Table 3.3. The $200 \text{ m}^2/\text{sec}$ value was selected to be sufficiently large that the filter would almost completely rely on the incoming measurements for the position estimation.

Assuming an average of eight satellites, the state vector for the filter was 16 states with the relative velocity and relative acceleration states (i.e., the full filter). In position-only mode, the state vector was reduced from 16 states to only 10 states. If the position-only filter still converged properly, and the results for the minimum indicator position error were similar to the full filter, then there would be no need for relative velocity or relative acceleration states in this application.

To test this theory, each data set was post-processed with the position-only filter. The position-only filter performed almost identically to the filter with the relative velocity and relative acceleration states included. The following items were compared to the real-time data: the time to the first fix of the true ambiguity set, the total number of unique ambiguity sets returned from the LAMBDA function, and the MRSE for the minimum indicator variable. If these items were similar between each of the data sets, then it would be safe to use the position-only filter in this application. Only flight test data sets were used for analysis, since that was the primary application for the algorithm.

Some of the post-processed data did not experience the same filter resets as during real-time, and thus the results from the real-time and post-processed runs were significantly different. Those data sets are denoted with an asterisk. Finally,

only the time to the fix for the first true ambiguity set is shown. If a filter reset occurred which caused the minimum indicator to reset, the next time to fix for the second true ambiguity set is not shown. This was primarily due to the fact that not all of the post-processed data sets experienced the same filter resets. Table 4.9 summarizes the results from each data run.

Based on the results from Table 4.9, there was no measurable difference between the two filter modes. Case F5.2 had 11 fewer unique ambiguity sets in the position-only mode for an unknown reason. Otherwise, the number of unique ambiguity sets returned from the LAMBDA function were always within two sets, and usually within one (except for Cases F2.1 and F3.2, for which the exact number of unique ambiguity set were unknown). The time to first fix was identical for every set. Since the true ambiguity set was found at the exact same time for every data set, the minimum indicator MRSE was necessarily the same as well (except for the data sets with filter resets). The results are not entirely surprising since both filters relied so heavily on the incoming measurements. The results from the position-only filter closely match those from the relative velocity and relative acceleration filter. Therefore, it is safe to use the position-only filter for future applications. The widelane example is discussed next.

Table 4.9: Velocity/Acceleration vs. Position-only

Data Set	Number of Unique Ambiguity Sets Returned from LAMBDA	Time to First Fix (minutes)	MRSE (meters)
F1.1 Velocity/Acceleration	29	4	0.030*
F1.1 Position-only	29	4	0.030*
F1.2 Velocity/Acceleration	41	4	0.085
F1.2 Position-only	27	4	0.069*
F2.1 Velocity/Acceleration	Unknown	4	0.062
F2.1 Position-only	27	4	0.033*
F2.2 Velocity/Acceleration	21	4	0.023
F2.2 Position-only	22	4	0.023
F3.1 Velocity/Acceleration	28	4.83 $\overline{3}$	0.137
F3.1 Position-only	30	4.83 $\overline{3}$	0.137
F3.2 Velocity/Acceleration	50+	4	0.153
F3.2 Position-only	50+	4	0.153
F4.1 Velocity/Acceleration	31	4.83 $\overline{3}$	0.243
F4.1 Position-only	30	4.83 $\overline{3}$	0.243
F5.1 Velocity/Acceleration	18	4	0.011
F5.1 Position-only	17	4	0.011
F5.2 Velocity/Acceleration	31	4	0.050
F5.2 Position-only	20	4	0.050
F6.1 Velocity/Acceleration	30	4	0.010
F6.1 Position-only	30	4	0.010
F6.2 Velocity/Acceleration	30	4	0.010
F6.2 Position-only	30	4	0.010
F6.3 Velocity/Acceleration	33	4	0.007
F6.3 Position-only	33	4	0.007
F6.4 Velocity/Acceleration	20	4	0.007
F6.4 Position-only	20	4	0.007
F7.1 Velocity/Acceleration	25	4	0.008
F7.1 Position-only	25	4	0.008
F7.2 Velocity/Acceleration	25	4	0.007
F7.2 Position-only	25	4	0.007
F8.1 Velocity/Acceleration	29	4	0.012
F8.1 Position-only	28	4	0.012
F8.2 Velocity/Acceleration	30	4	0.016
F8.2 Position-only	30	4	0.016

*Note: Post-processed data which did not experience same filter reset(s) as real-time data

4.5.3 Widelane Observables Comparison. All the results presented in thesis thus far were L1 only phase observables, due to the Javad® JNS100 receiver. Assuming the receivers were dual channel capable, a new measurement observable called widelane (WL) can be formed. The formation of the widelane observable was defined in Section 2.2.5.

The algorithm in this thesis was coded for use with either L1 or WL observables. A set of WL measurements were constructed from Ashtech GPS Z-Surveyor data taken on 28 May 2004 at AFIT. The data set was static with a known baseline distance. The receivers were approximately 1450 meters apart. The WL data run was set up slightly differently from what was done in the previous examples. There were two major differences. The first was the filter convergence time before the floating ambiguities were sent to the LAMBDA function. Second, the number of SOS residual points collected before performing the least squares fit calculation was significantly reduced. As described in Sections 3.5 and 3.7.3, the filter was given a two minute convergence time before the floating filter ambiguity estimates were sent to the LAMBDA function, and two minutes worth of SOS residuals were collected before the least squares fit calculation was performed. For the WL example, the filter was not given any convergence time, and the least squares fit was performed after 20 seconds. Twenty seconds resulted in three SOS residuals data points to perform the least squares fit, the minimum required.

Figure 4.40 and Figure 4.41 show the WL and L1 SOS residuals respectively. The scale on the L1 plot is expanded more than that of the WL plot to show the true ambiguity set better; at 0.09 meters on the y-axis, the true ambiguity set was unidentifiable in the L1 plot. The difference between the two plots is dramatic. In the WL plot, the true ambiguity set is clearly identifiable, even from the very first epoch. The L1 example is another story. There were multiple erroneous ambiguity sets' residuals intermingled with the true set. Also, there were at least seven different local minima where an erroneous ambiguity set's residuals were below the true set's

residuals. The magnifying glass analogy discussed earlier is dramatically evident in the WL plot compared with the L1 plot.

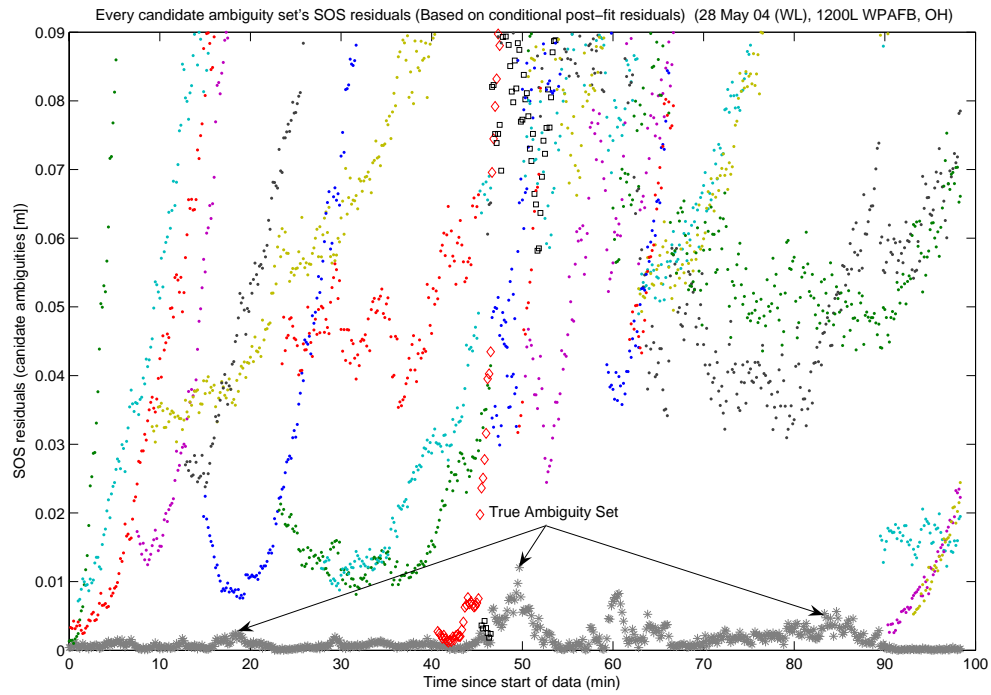


Figure 4.40: WL SOS Residuals

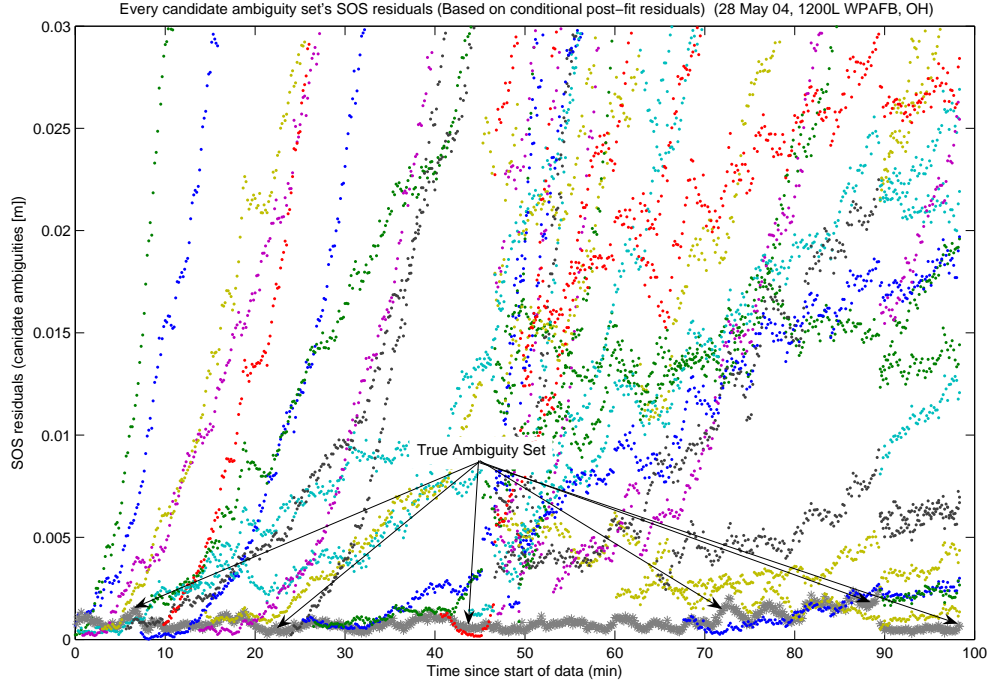


Figure 4.41: L1 SOS Residuals

An expanded plot of the WL SOS residuals is shown in Figure 4.42. The true ambiguity set's residuals are lower than all the erroneous ambiguity sets' residuals during the entire data run, even on the first epoch.

Next, the conditional probability plots for the WL and L1 observables are shown in Figures 4.43 and 4.44 respectively. The true ambiguity set immediately absorbed all of the probability in the WL example, and never dipped below 93 percent. However, in the L1 example, the true ambiguity set did not absorb all of the probability until 20 minutes, and then dropped three more times, until 90 minutes when it absorbed all the probability for the final time.

Finally, Table 4.10 summarizes the results from each of the data runs. The WL run “found” the true ambiguity set a full eight minutes before the L1 run. The L1 minimum indicator chose two erroneous ambiguity sets before selecting the true set at 8.33 minutes. Interestingly, both runs had about the same number of unique ambiguity sets returned from the LAMBDA function. Since many current receivers are dual channel capable and only small changes to the filter are required to use

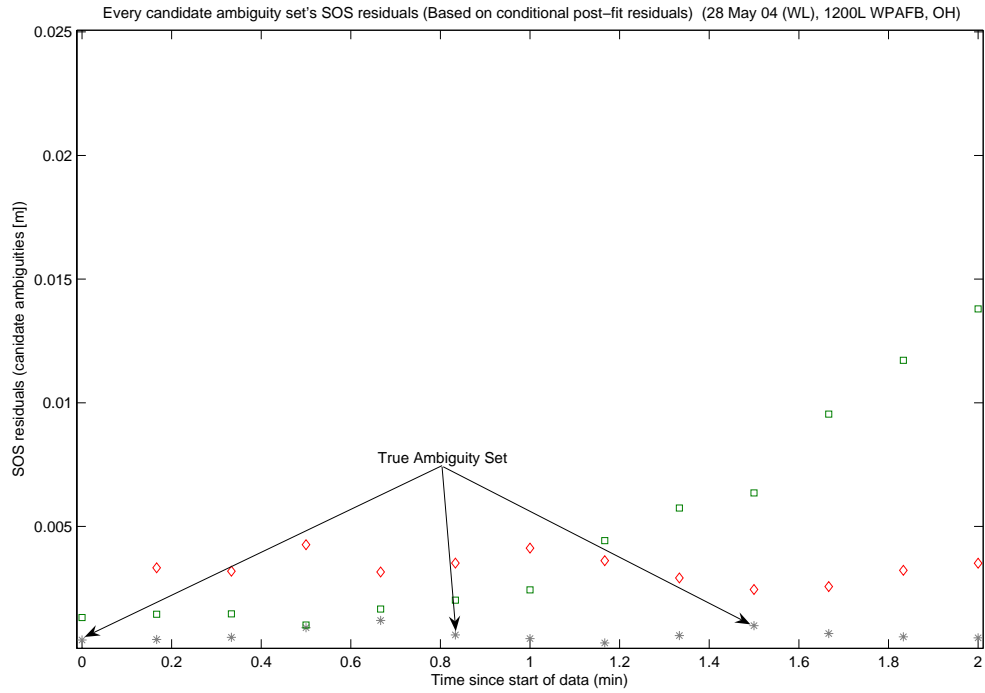


Figure 4.42: WL SOS Residuals Expanded

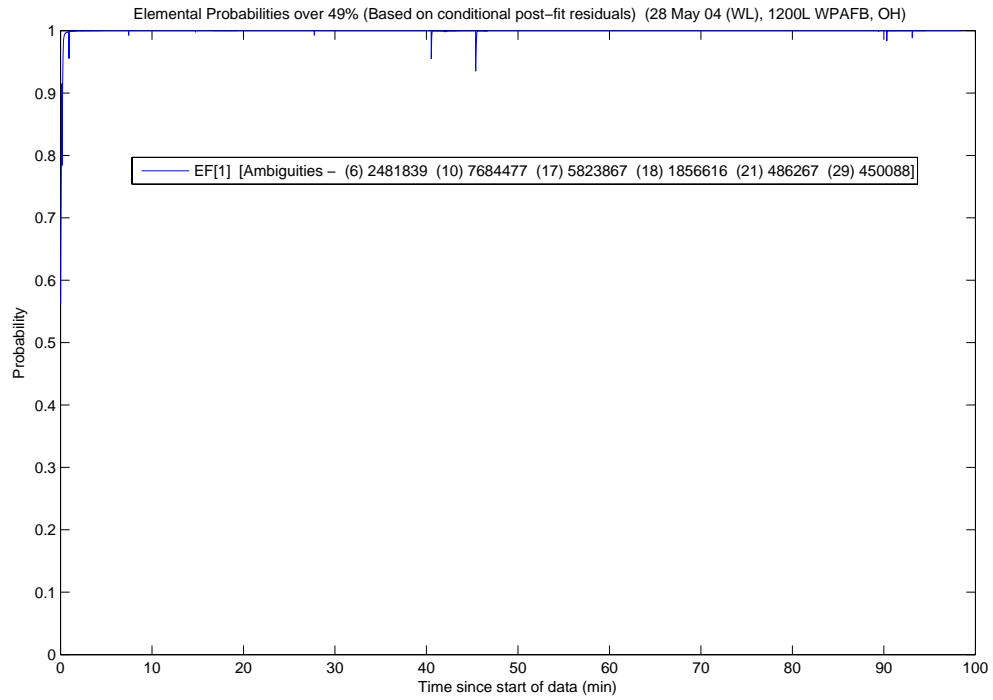


Figure 4.43: WL Conditional Probabilities for Selected MMAE Elemental Filters

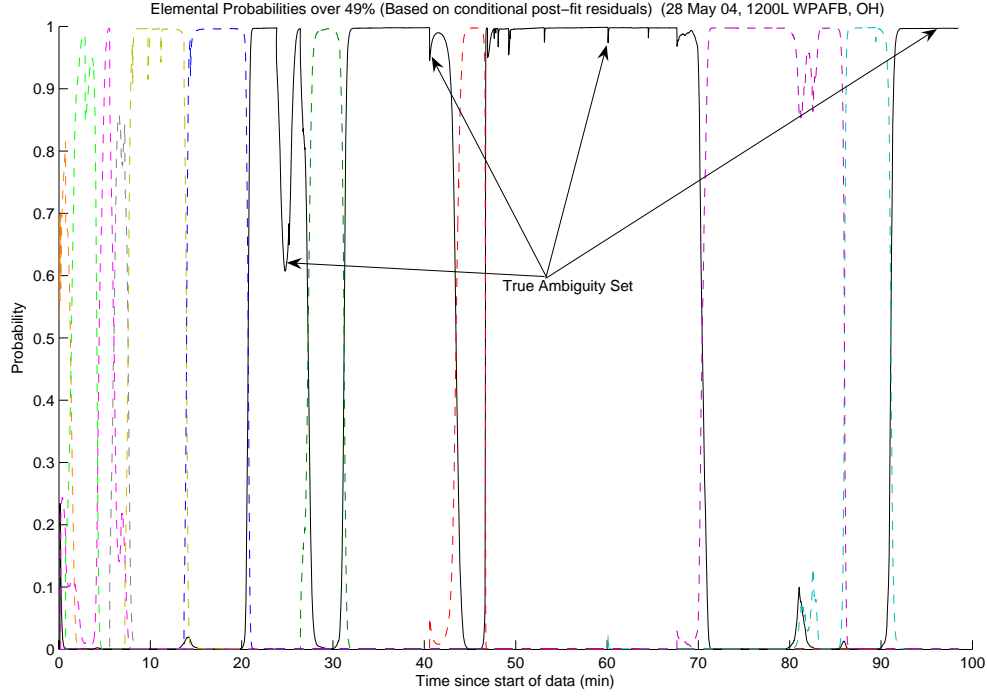


Figure 4.44: L1 Conditional Probabilities for Selected MMAE Elemental Filters

widelane observables, the benefits from widelaning far outweigh the extra cost of a dual channel receiver. The only caveat is the receiver must have the ability to output the data from both frequencies at a sufficient rate (greater than or equal to 20 Hz) for use in applications similar to the one tested in this thesis. The last column in Table 4.10 shows the MRSE for the time period when both measurement observables had the true ambiguity set selected. As expected, the position error on the L1 observables is less than the WL observables. A dynamic T_a and σ_a^2 comparison is shown next.

Table 4.10: Widelane and L1 Observables Summary

Data Set	Number of Unique Ambiguity Sets Returned from LAMBDA	Time to First Fix (minutes)	MRSE (meters)	MRSE* (meters)
Widelane	27	0.33	0.051	0.051
L1	30	8.33	0.272	0.025

*Note: This final MRSE was calculated from 8.33 minutes until the end of the data run, when both measurement observables had the true ambiguity set selected.

4.5.4 *Dynamic T_a and σ_a^2 Comparison.* The final special case is a demonstration of dynamic T_a and σ_a^2 in the FOGMA model described in Section 3.4.1. The dynamic T_a and σ_a^2 should make the floating filter's 1σ estimate of the relative velocity and relative acceleration more accurate. The dynamic T_a and σ_a^2 example was post-processed, so the errors on the relative velocity and relative acceleration states do not exactly match the real-time data. Figure 4.45 shows the real-time relative velocity error plot from 18 April 2005. The floating filter's 1σ estimate of the relative velocity errors remained constant throughout the run with a static T_a and σ_a^2 . However, when T_a and σ_a^2 are allowed to change dynamically based on the total relative acceleration vector, the filter more accurately models the relative velocity errors. Figure 4.46 shows the same data set with dynamic T_a and σ_a^2 . Although not perfect, the dynamic T_a and σ_a^2 clearly modeled the velocity errors much better than the static case.

Next, the relative acceleration errors are shown. The plots are similar to the relative velocity error plots. Once more, the dynamic T_a and σ_a^2 modeled the relative acceleration errors much more closely than the static case. Figure 4.47 shows the static and Figure 4.48 shows the dynamic case.

Assuming accurate estimates of the error on the relative velocity and relative acceleration states are needed, the dynamic filter provides a quick, easy solution. Proper tuning of the T_a and σ_a^2 values, based on system dynamics, could provide an even more accurate estimate of the estimated errors than the results shown in this thesis. Only minimal changes to the basic filter are required to implement dynamic T_a and σ_a^2 .

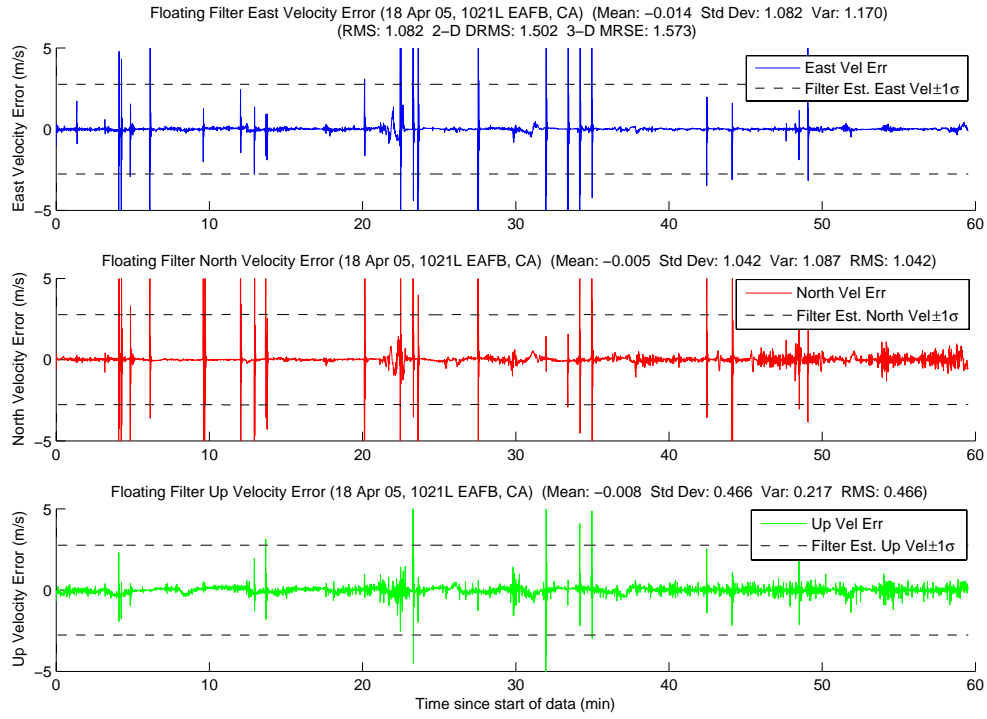


Figure 4.45: Static T_a and σ_a^2 Floating Filter Velocity Errors

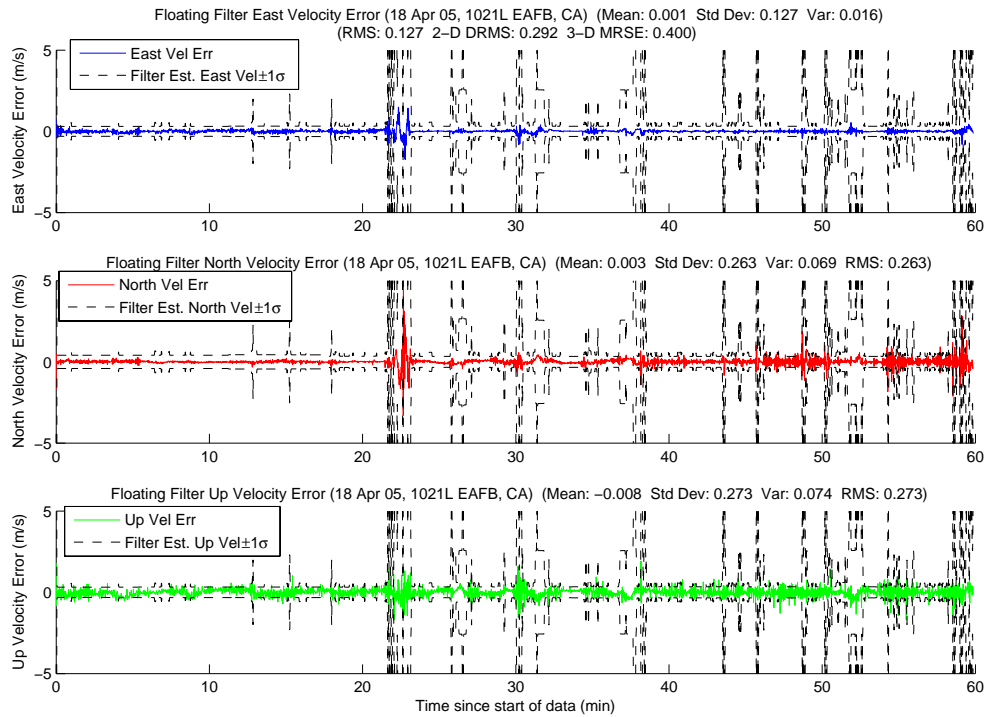


Figure 4.46: Dynamic T_a and σ_a^2 Floating Filter Velocity Errors

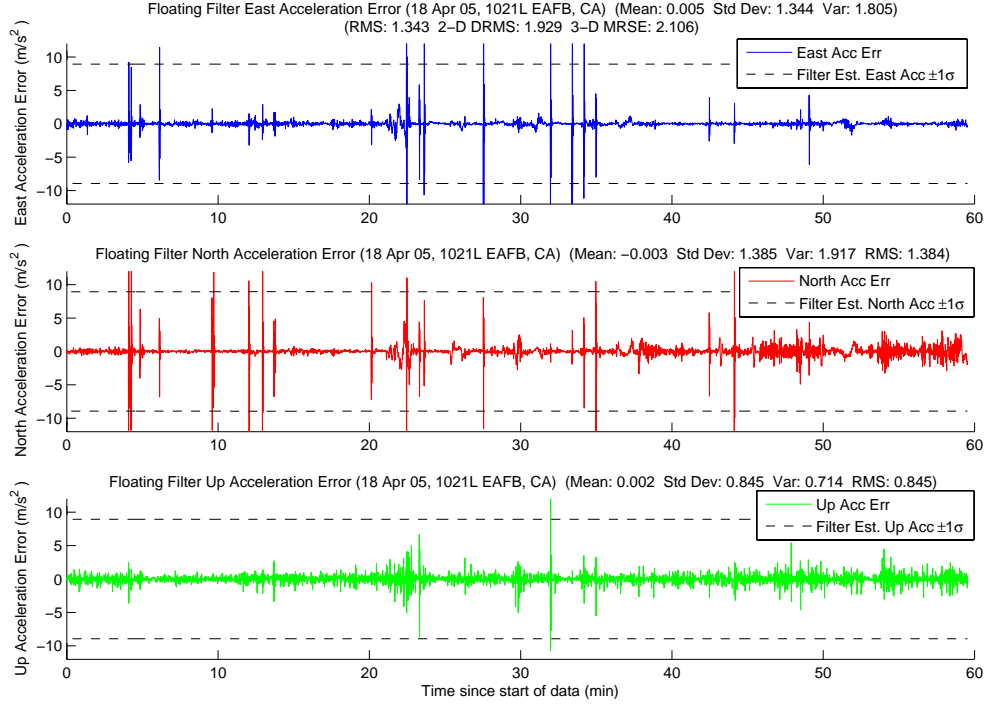


Figure 4.47: Static T_a and σ_a^2 Floating Filter Acceleration Errors

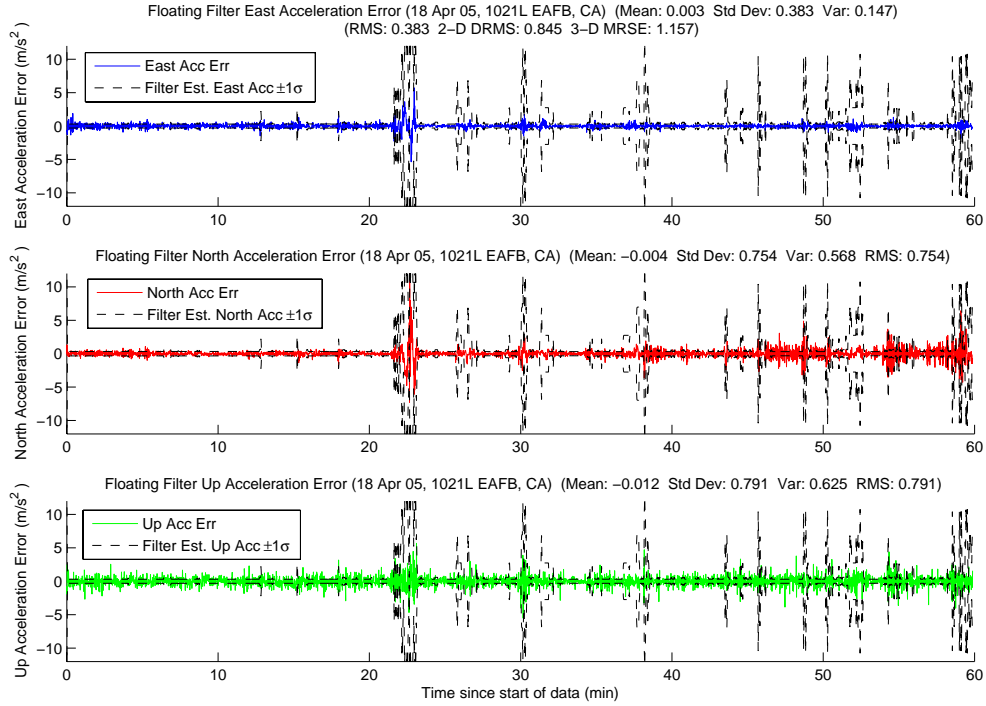


Figure 4.48: Dynamic T_a and σ_a^2 Floating Filter Acceleration Errors

4.5.5 Special Case Summary. Four special cases were presented: “pre-fit” and “post-fit” residuals, the filter in position-only mode, widelane measurement observables example, and the floating filter with the dynamic T_a and σ_a^2 option turned on. The equivalence of “pre-fit” and “post-fit” residuals in the probability calculation was verified. Also, the special cases showed possible options to reduce the computational load, increase the floating filter’s ability to resolve the true ambiguities, and improve the floating filter’s 1σ estimate of the relative velocity and relative acceleration states.

4.6 Summary

Chapter IV presented the results and analysis from ground and flight testing. Included in the results were a discussion of the test methodology and data set descriptions. Next, the results from two ground tests were presented, static and dynamic. Following the ground tests, specific flight tests were analyzed. Finally, special cases such as, “pre-fit” and “post-fit” residual comparison, the floating filter in position-only mode, widelane measurements, and dynamic T_a and σ_a^2 were discussed. All objectives were met, and the system achieved the centimeter-level requirement in all but four out of seventeen data runs.

V. Conclusions and Recommendations

5.1 Overview

The goal of this thesis was to develop and test a high-speed, precise DGPS system for use with UAVs. To get the required precision for close formation flight, a new ambiguity resolution technique was developed, based on special characteristics of the carrier-phase residuals. In Chapter II, the theory behind design of the system was presented including: EKF, MMAEs, GPS theory, and the nature of the phase residuals. Chapter III described the hardware configuration, overall algorithm, and system models used in implementation of the theory from Chapter II. Finally, the results and analysis from ground and flight testing were presented in Chapter IV. The end result was the first-ever successful in-flight demonstration of automated close formation flight, culminating in over 11 hours with a mean radial spherical error of 3.3 centimeters (0.108 feet). With this type of accuracy, automated aerial refueling is now possible for UAVs and tankers equipped with DGPS, a datalink, and software developed in this thesis. A concise summary of the conclusions and recommendations follows next.

5.2 Conclusions

- **The high-rate precise position system as a proof-of-concept was validated.** Although the system was not without problems, it demonstrated that a high-rate precise DGPS solution is adequate for derivative-type flight controllers. Automated flight position changes between two aircraft were safely conducted over 10 hours of flight testing during the October 2005 flight tests. By the end of the flight testing, the pilots were able to “flip-the-switch” and watch the system safely operate without any outside intervention.
- **Consistent centimeter-level positioning was attained.** Including ground and flight testing, there were 19 data sets for a total of approximately 16 hours of testing. Both ground tests achieved centimeter-level positioning. However, the primary interest of this thesis was the flight testing, specifically when the

aircraft were within 500 feet. During 16 of the 17 flight tests, the aircraft were within 500 feet. On the one flight in which the aircraft did not break the 500 foot barrier, the minimum indicator was within the centimeter-level requirement during the entire run.

The times when the aircraft were within 500 feet resulted in approximately 11 hours of data. During the 11 hours, there was only one 40 second period when the minimum indicator selected an erroneous ambiguity set due to datalink connection problems. Overall, the 3-D MRSE for the entire 11 hours was 3.3 centimeters.

- **The algorithm was robust enough to handle most potential issues.**

During the 16 hours of testing, numerous issues arose which needed to be handled properly. GPS satellites rose and set over the horizon similar to the sun, necessitating the addition and removal of satellites during data runs. Code was put into place to deal with both cases. The algorithm correctly handled every satellite addition and removal.

The reference satellite was chosen based on the elevation angle above the horizon. More highly elevated satellites were not as prone to masking and multipath. In addition, the satellite signal traveled through less atmosphere so the measurement errors were usually smaller. However, since the satellites were in constant motion, there was the potential for the maximally elevated satellite to change during a data run. Instead of resetting the filter and losing all of the information garnered up to that point, a transformation of the state vector and covariance matrix was applied. The algorithm correctly handled all of the maximally elevated satellite changes.

Since the phase measurements were being used for the final precise relative position, it was imperative that any potential cycle slips were detected and resolved. The algorithm set conservative limits for the two methods employed for

cycle slip detection. As a result, there were a few erroneous cycle slip detections, but every real cycle slip was identified and corrected.

Finally, datalink errors needed to be addressed. Generally, the algorithm handled small datalink transmission errors, such as CRC failures and lost messages. However, due to the conservative values in the cycle slip detection methods, when long strings of messages were “lost”, the algorithm would erroneously detect cycle slips on all of the satellites, resulting in a system reset. This problem usually appeared when the aircraft were beyond 500 feet from each other. There was one instance, mentioned in Chapter IV, in which this problem occurred when the aircraft were within 500 feet.

- **Occasionally, a timing anomaly resulted in the skipping of a single precise relative position message each second, as seen by the flight controller.** The algorithm received raw GPS measurements from the Javad® JNS100 receiver at a 20 Hertz rate (i.e., a message every 50 milliseconds). The precise relative position output should have been at that rate as well. Unfortunately, it would generally stagger between 31 and 62 milliseconds. Occasionally, the delay was 93 milliseconds or higher, but was always a multiple of 31 milliseconds. The exact cause of the 31 millisecond stagger is not entirely known, but is most likely due to context switching from the multiple algorithm threads running on one processor. To investigate the one second periodic nature of the problem further, the inter-message delay between 20 successive messages was summed. The total was usually one second as expected, but occasionally was 1.031 seconds. The extra 31 milliseconds corresponded to exactly one “standard” message interval. Finally, the times when the flight controller skipped a relative position message correlated to the times there was the additional 31 milliseconds.

The missing messages initially caused large errors in the flight controller, and the in-flight variable stability system would disengage. To overcome this problem, the flight controller would estimate the relative position based on the

previous two epochs when a message was “lost”. This work-around eliminated any deleterious effects of the missed data epochs.

- **The newly developed minimum indicator method provided a more stable and precise relative position output when compared against the MMAE approach.** The average minimum indicator 3-D MRSE was an order of magnitude less than the average MMAE 3-D MRSE over the course of all 17 flight tests (0.095 meters vs. 0.215 meters). There was only one case for which the MMAE 3-D MRSE was lower than the minimum indicator 3-D MRSE during the 17 flight tests. This occurred because the minimum indicator chose an erroneous ambiguity set for the first minute of the data run while the aircraft were rejoining. After the aircraft joined, the minimum indicator remained within the centimeter-level requirement for the remainder of the run.
- **The number of visible satellites and their associated geometry plays an extremely important role in nature of the SOS residuals.** As evidenced by Cases 5.1 and 5.2, the determination of the true ambiguity set becomes much easier with more visible satellites. With each new measurement, it becomes more difficult for an erroneous ambiguity set to fit the measurements better than the true set. The residual plot from Case 5.1, for which there are 12 visible satellites, looks very similar to the widelane example.
- **“Pre-fit” and “post-fit” residuals result in the same MMAE conditional probability calculation.** The theory proposed by Ormsby [44] was validated in this thesis. As seen in Chapter IV, there was almost no difference between the two calculations even though the “pre-fit” and “post-fit” calculations came from different methods.
- **A simplified dynamics model can be used without a noticeable degradation in performance.** When there were no filter resets, the position-only filter performed identically to the full filter implementation. If estimates of the relative velocity and relative acceleration are not needed, the processing require-

ments can be significantly reduced by using the position-only filter. This allows processing for additional enhancements or future growth.

- **Widelane observables offer almost instantaneous ambiguity resolution at the relative ranges in this thesis.** The benefit of widelane observables was clearly evident from the results in Chapter IV. Even when the baseline relative distance was almost 1.5 kilometers, the true ambiguity set was the top set returned from the LAMBDA function on the first epoch. Also, the minimum indicator selected the true ambiguity set after the first three SOS residual calculation epochs (20 seconds) and remained “locked” for the duration of the run. Using the L1-only measurements, the algorithm took eight minutes longer to determine the true ambiguity set, a 2400 percent increase! The main drawback to the widelane observables is the increased position error. However, the increase in the error was only 2.5 centimeters in this example. The incredible increase in performance far outweighs the small increase in error.
- **A dynamic correlation time constant, T_a , and dynamic mean squared value, σ_a^2 , for the FOGMA model significantly enhances the floating filter’s 1σ estimates of the relative velocity and relative acceleration states.** As expected, the floating filter more closely modeled the relative velocity and acceleration errors with dynamic T_a and σ_a^2 values as seen in Figures 4.46 and 4.48. If accurate estimates of the relative velocity and relative acceleration states are required, dynamic T_a and σ_a^2 offer a simple and efficient method. However, there is an increase in the computational requirements since the dynamics matrix, \mathbf{F} , dynamics noise matrix, \mathbf{Q} , and the state transition matrix, Φ , must be recomputed each time T_a and σ_a^2 change.

5.3 Recommendations

The system proved to be robust and reliable. However, some issues were found during testing which warrant further research. The recommendations listed below summarize the more important issues.

- **Determine exact cause of timing problem.** The biggest problem encountered during testing was the occasional timing issue resulting from extended inter-message delays. Higher speed precise relative position output will not be possible until this issue is resolved. If it is due to the processor, then new dual core processors might offer a potential fix. Also, streamlining the number of program threads may help improve the algorithm’s performance.
- **Explore other variations on the minimum indicator equation.** The minimum indicator added each element in the numerator of the equation. However, multiplication might also be a viable option, or some other functional form to include using weighting. A new form could potentially improve the performance of the minimum indicator, especially to system changes such as cycle slips.
- **Use the magnitude of the SOS residuals as an additional discriminator of the true ambiguity set.** Since the algorithm divided by the number of epochs an ambiguity set was “alive” and the number of epochs it was returned as the top set from LAMBDA, the minimum indicator value for the true ambiguity set continually decreased throughout a data run. Unfortunately, if a cycle slip occurred, it could take a significant amount of time for the minimum indicator value of the previously true ambiguity set to increase, especially if it has been “alive” for an extended time. This prevented the minimum indicator from the selecting the true ambiguity set in a timely fashion, assuming it was in the candidate ambiguity bank. Case G2 was a good example of this phenomenon.

After examining the SOS residuals plot for each data run, the true ambiguity set’s residuals exceeded three millimeters in one case, and only for a few epochs. An SOS residual ceiling could be imposed, based on vehicle dynamics and slant range between the aircraft, which would prevent the minimum indicator from selecting a candidate ambiguity set as true if its SOS residuals exceeded the ceiling.

- **Use the position-only EKF (i.e., the simplified dynamic model).** Based on the results from Chapter IV, there is no need to run with relative velocity and relative acceleration states in the EKF. The main output from the algorithm is the high-speed precise relative position estimate, not estimates of the relative velocity and relative acceleration. The reduction in the EKF states will make the algorithm quicker and more efficient.
- **Acquire a high-speed dual frequency receiver.** A dual frequency receiver will accomplish two things. First, it will make determination of the true ambiguity set much easier, almost instantaneous. Next, it will allow for more efficient and precise methods to check and correct cycle slips [14].
- **Increase the MMAE blended solution probability cutoff.** Any ambiguity set with a conditional probability at or below one percent was excluded from the MMAE Bayesian position estimation. Although this helped make the MMAE more accurate, a much larger cutoff, such as 30 to 40 percent, could increase the MMAE accuracy even more.
- **Improve the post-processing algorithm to match the real-time algorithm more closely.** One drawback of the algorithm was the inability to make the post-processed data return exactly the same results as the real-time data. By placing a time stamp on each message during real-time execution, the post-processed data could reflect the real-time data more accurately.

Appendix A. Data Set Plots

A.1 Ground Tests

Two ground tests are discussed in this section. The first was a static case taken at the AFIT navigation laboratory on 25 March 2005. Next, a dynamic driving case is shown. The dynamic case shows some of the errors discussed in Chapter III.

A.1.1 Case G1: Static Ground Test. Case G1 is the only static data set discussed, other than the widelane observables. The algorithm easily handled the static case. Figures A.1, A.2, and A.3 display the floating filter’s relative position, velocity, and acceleration errors, respectively. The dashed lines represent the floating filter’s 1σ estimation.

The title of each subplot contains the associated statistics for that plot. Additionally, the distance root mean square (DRMS) and mean radial spherical error (MRSE) are listed in the first (East) subgraph title. The DRMS can be thought of as a horizontal 2-dimensional indicator of error, and the MRSE as a spherical 3-dimensional indicator of error. These two values define system performance better than the mean and standard deviation alone. The MRSE was used to determine the algorithm accuracy for each data run. This value was then compared against the centimeter-level requirement.

After an initial convergence period of approximately 10 minutes, the floating filter “found” the correct relative position. As expected, the floating filter was very accurate for a static case with an extremely short baseline (1 meter). Based on the float filter’s 1σ estimation, the floating filter required more tuning for this particular case. However, the floating filter was designed to work with a variety of vehicle dynamics and therefore was not perfectly tuned for the static case. Finally, the floating filter was tuned to give the best possible outcome from the minimum indicator and the best MMAE weighted position estimate. In order to do this the LAMBDA algorithm must return the true ambiguity set. The floating filter was only a tool to help the LAMBDA function. Ensuring the LAMBDA algorithm returned the true

ambiguity set in the quickest possible time was the top priority, even at the expense of not tuning the floating filter for a theoretically accurate covariance matrix.

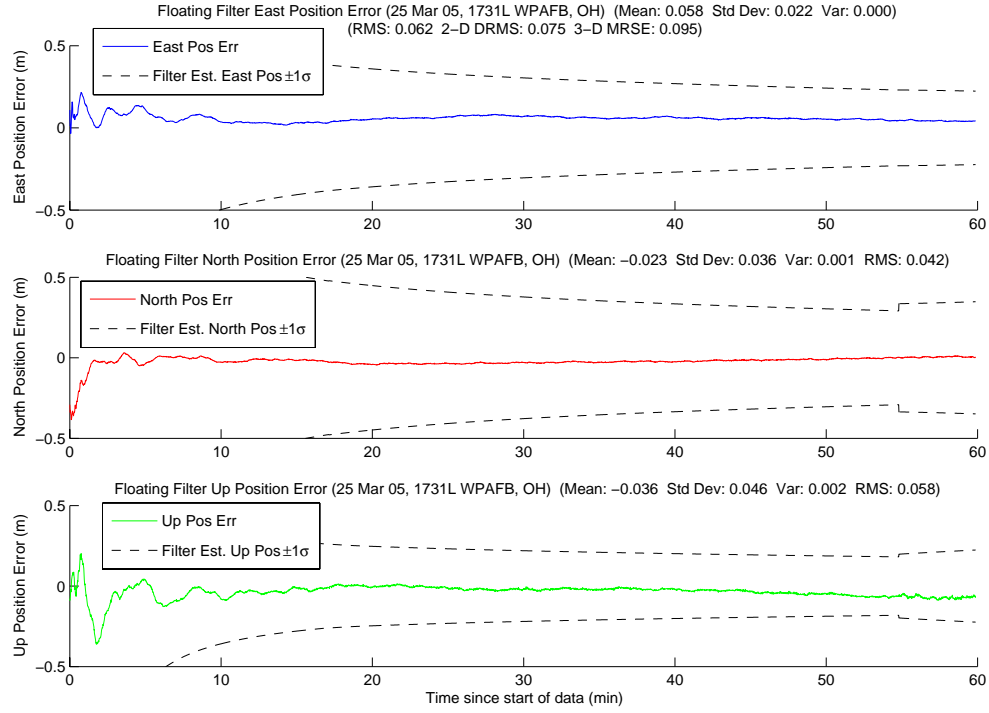


Figure A.1: Case G1: Floating Filter Position Errors

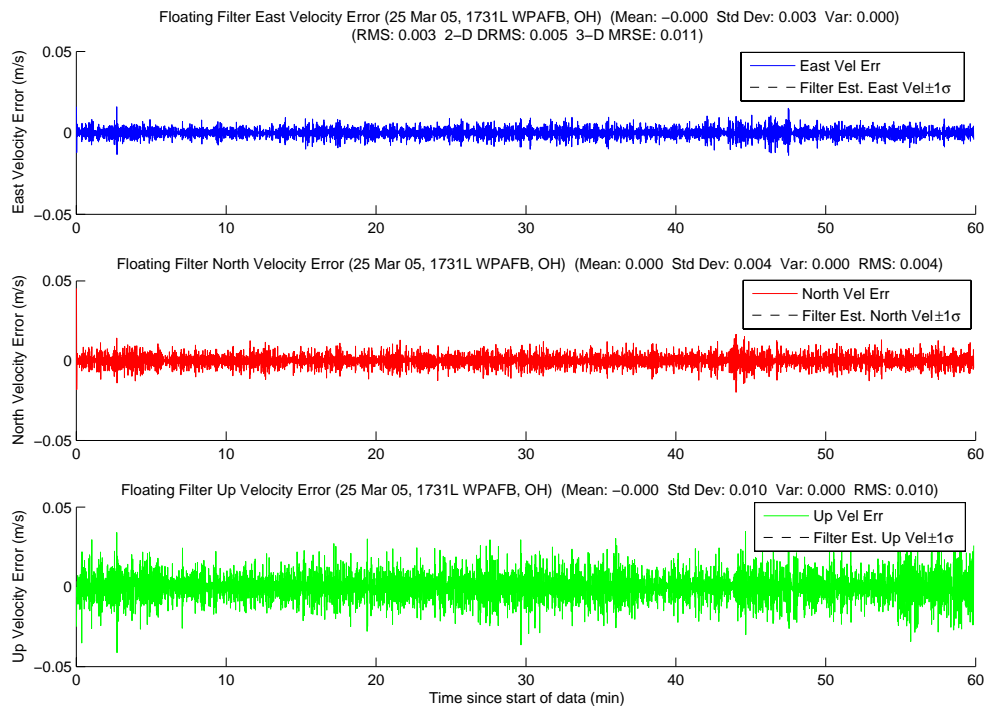


Figure A.2: Case G1: Floating Filter Velocity Errors

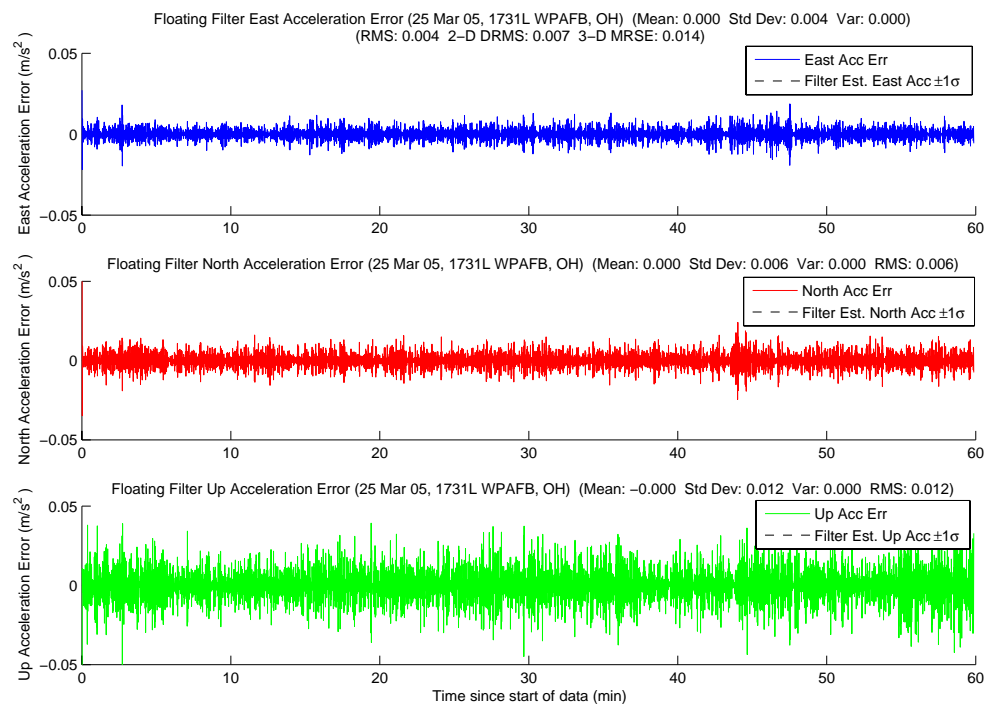


Figure A.3: Case G1: Floating Filter Acceleration Errors

Next, each satellite's measurement data is shown. Satellite 26 data is in Figure A.4. A small amount of code multipath can be seen in the cyclic nature of the pseudorange residuals. Also, the decrease in the 1σ for the ambiguity estimate at 18.5 minutes was due to a change in the reference satellite from pseudo-random noise PRN 18 to PRN 21. The rest of the satellite's plots follow in Figures A.5, A.6, A.7, A.8, A.9, A.10, and A.11. Figure A.12 shows the number of visible satellites throughout the data run.

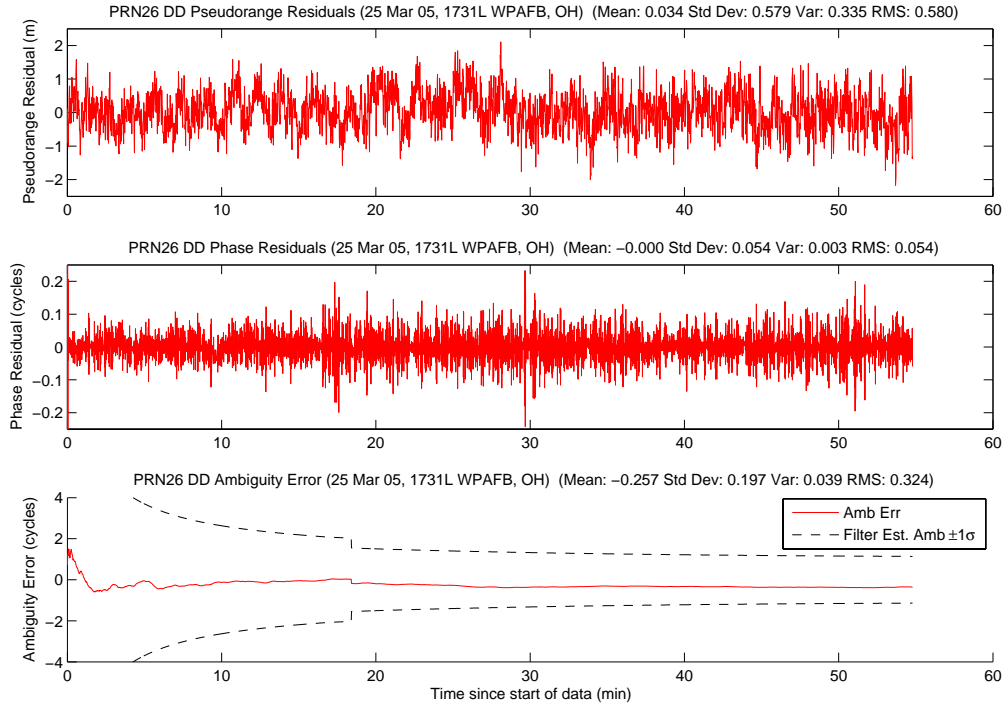


Figure A.4: Case G1: Satellite 26 Measurements

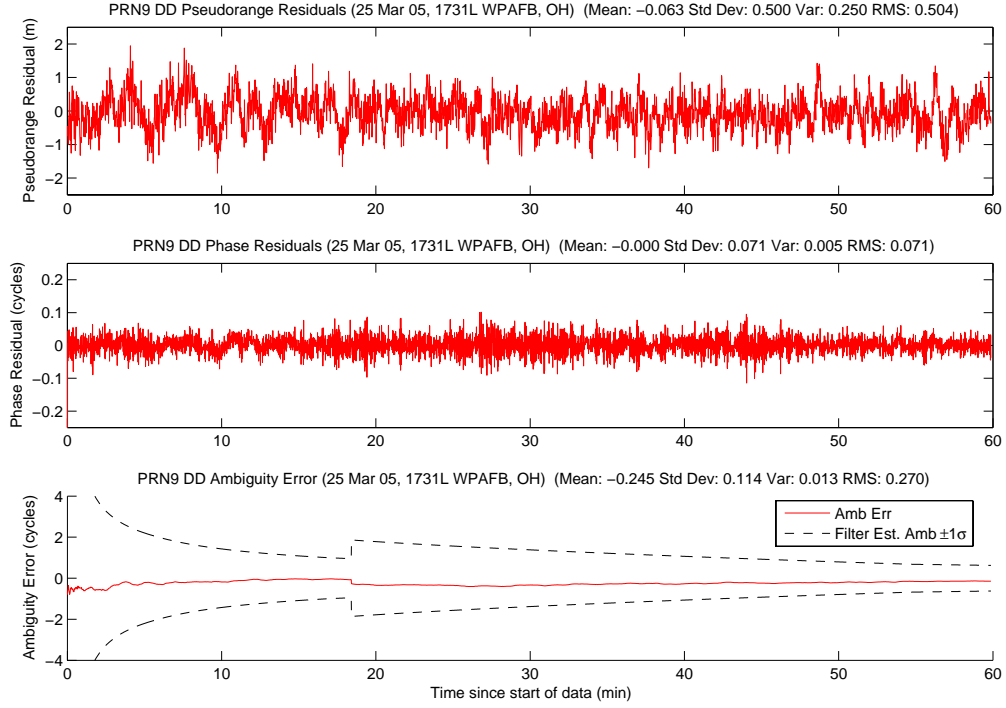


Figure A.5: Case G1: Satellite 9 Measurements

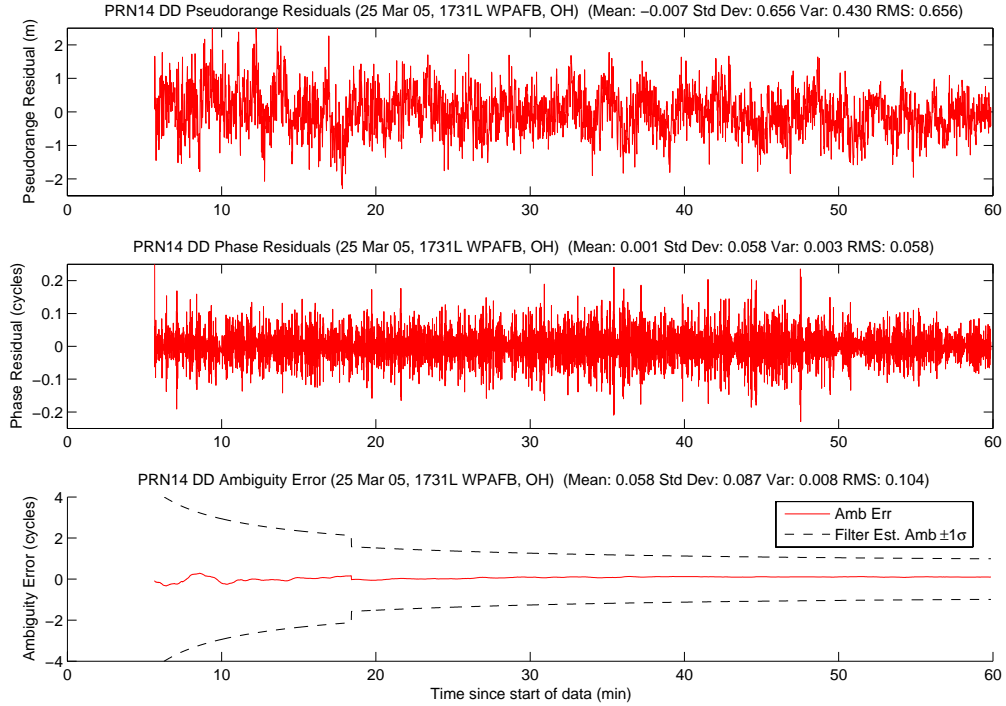


Figure A.6: Case G1: Satellite 14 Measurements

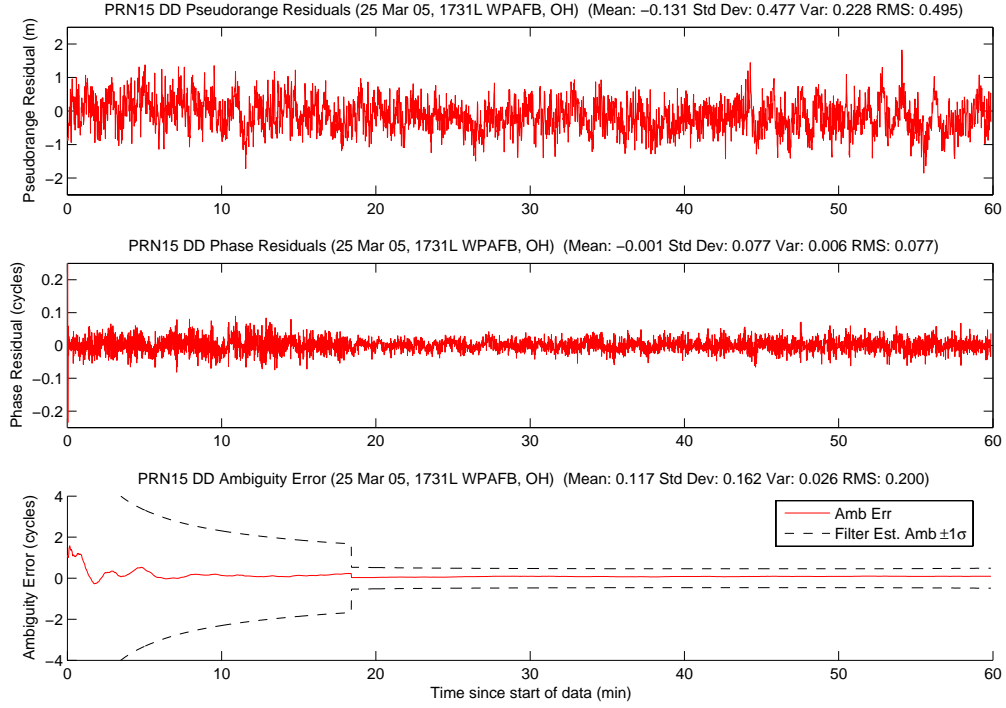


Figure A.7: Case G1: Satellite 15 Measurements

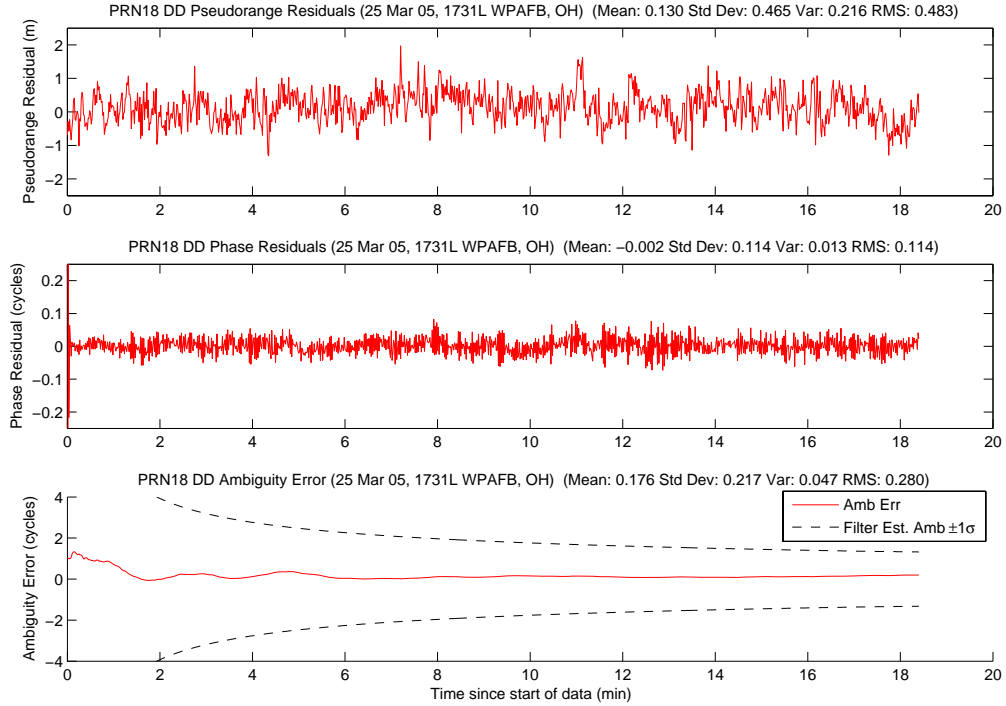


Figure A.8: Case G1: Satellite 18 Measurements

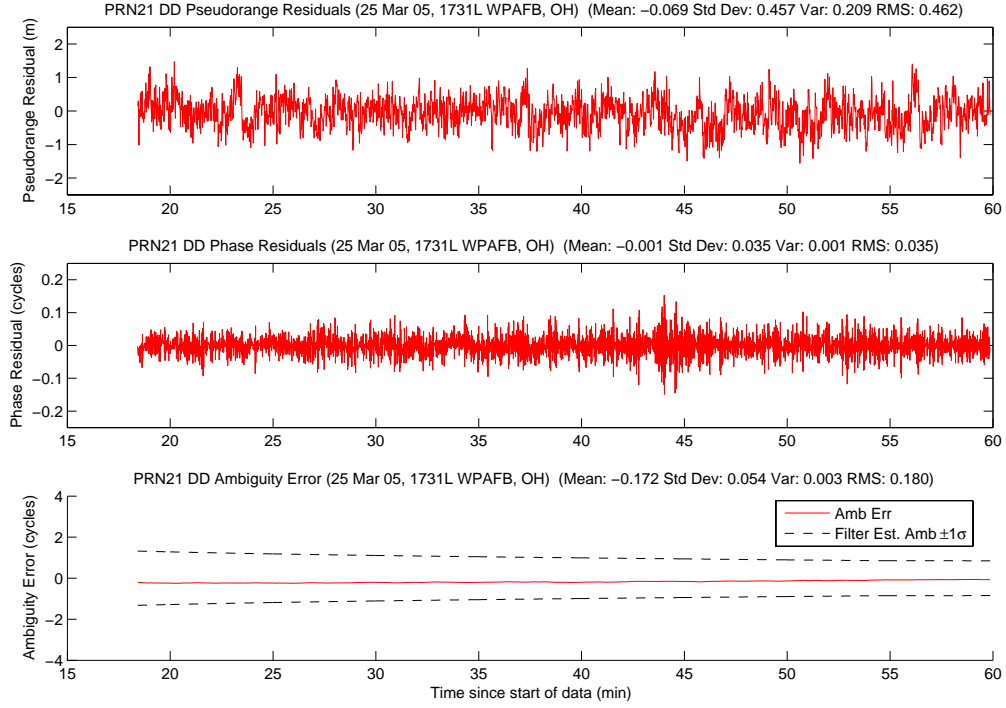


Figure A.9: Case G1: Satellite 21 Measurements

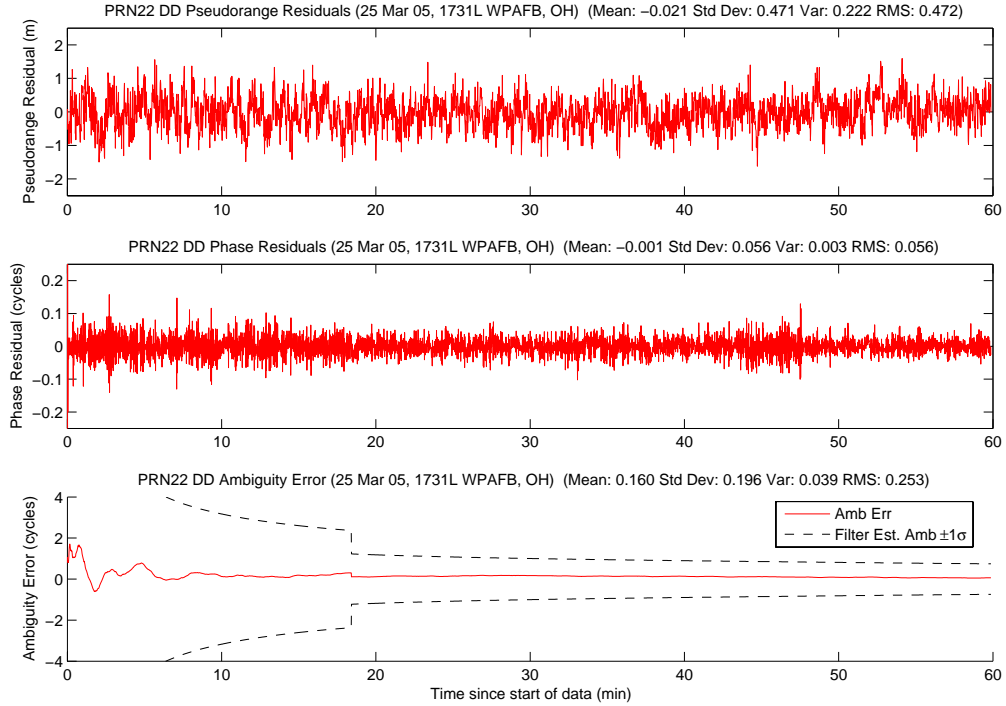


Figure A.10: Case G1: Satellite 22 Measurements

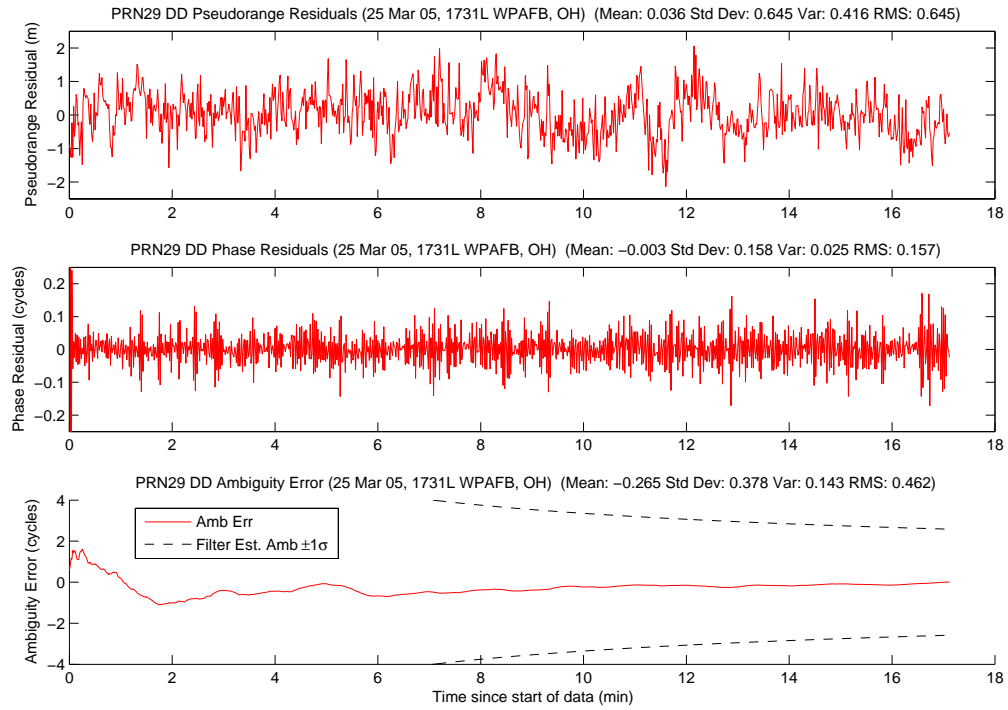


Figure A.11: Case G1: Satellite 29 Measurements

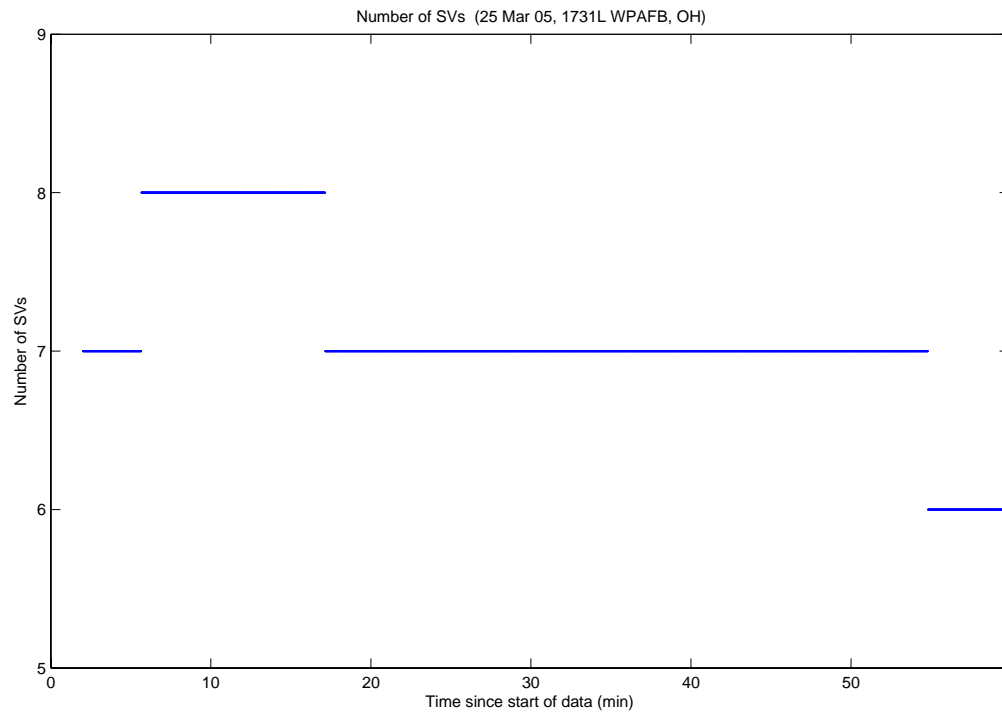


Figure A.12: Case G1: Visible Satellites

After examining the output files from the data run, there were a total of 21 unique ambiguity sets returned from the LAMBDA algorithm. Since there was a two minute convergence time for the floating filter prior to sending the floating filter ambiguity estimates and their associated covariance to the LAMBDA algorithm, and a two minute wait time prior to doing the minimum indicator calculations, the earliest the correct ambiguity set could be discovered was four minutes into a data run. In the first pair of candidates returned from LAMBDA, the true set was the second candidate set. The top candidate was later culled from the ambiguity bank as the floating filter converged to the correct relative position. This was the only candidate pruned during the data run. The minimum indicator initially selected the wrong ambiguity set at four minutes. However, at the next least squares fit epoch, ten second later, the minimum indicator correctly selected the true ambiguity set. The minimum indicator did not change throughout the rest of the data run. Also, the true ambiguity set was returned from the LAMBDA function as the top set, 3411 out of 3600 epochs (94.8 percent). Finally, the highest ratio reached by the true ambiguity set was 31.8.

An examination of the SOS residuals and conditional probability plots reveals more interesting information. All the candidate ambiguity set's SOS residuals are shown in Figure A.13 and the probabilities in Figure A.14. The EF[#] in the legend of the probability plot shows which elemental filter (EF) was associated with the ambiguity set displayed beside it. A slash in the EF number section indicates a lower EF candidate set was pruned and the elemental filter's new number is shown. In this example, the number one EF was pruned. Therefore, all the elemental filters moved down one position. The parenthesis in the legend are the PRN numbers with their associated double-difference ambiguity. The true ambiguity set is explicitly identified with an arrow when required. To reduce clutter, only candidate sets which rose above 49 percent are displayed.

One interesting aspect of Figure A.14 is when the MMAE “found” the true ambiguity set. This occurred at four minutes, approximately the same time the

minimum indicator variable determined the correct set. Since each technique was independent of the other, this did not always occur.

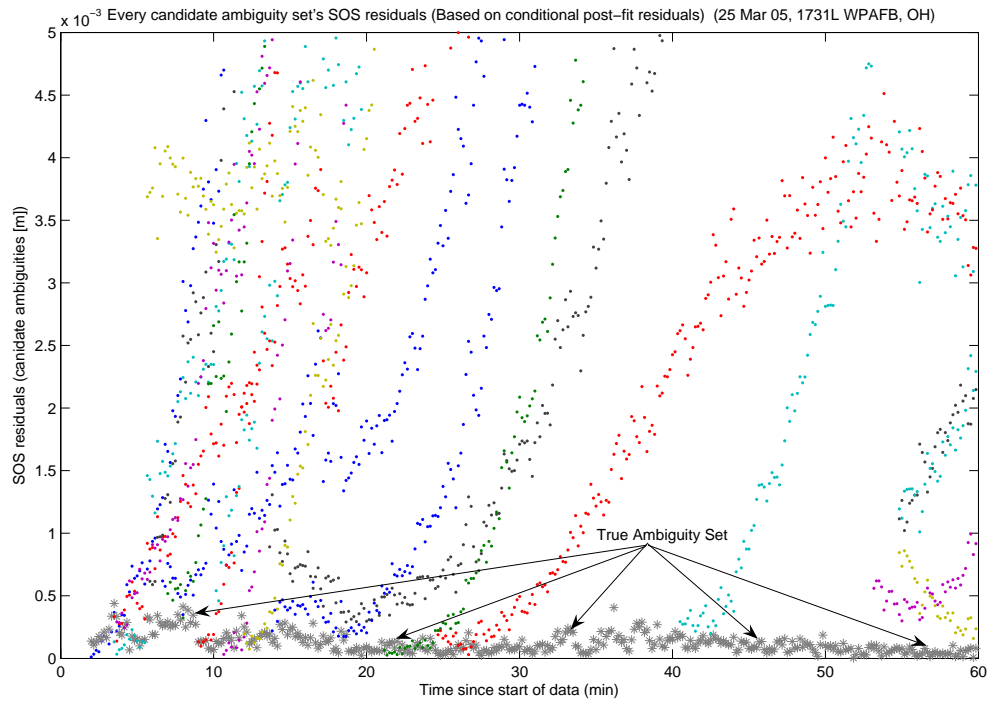


Figure A.13: Case G1: SOS Residuals

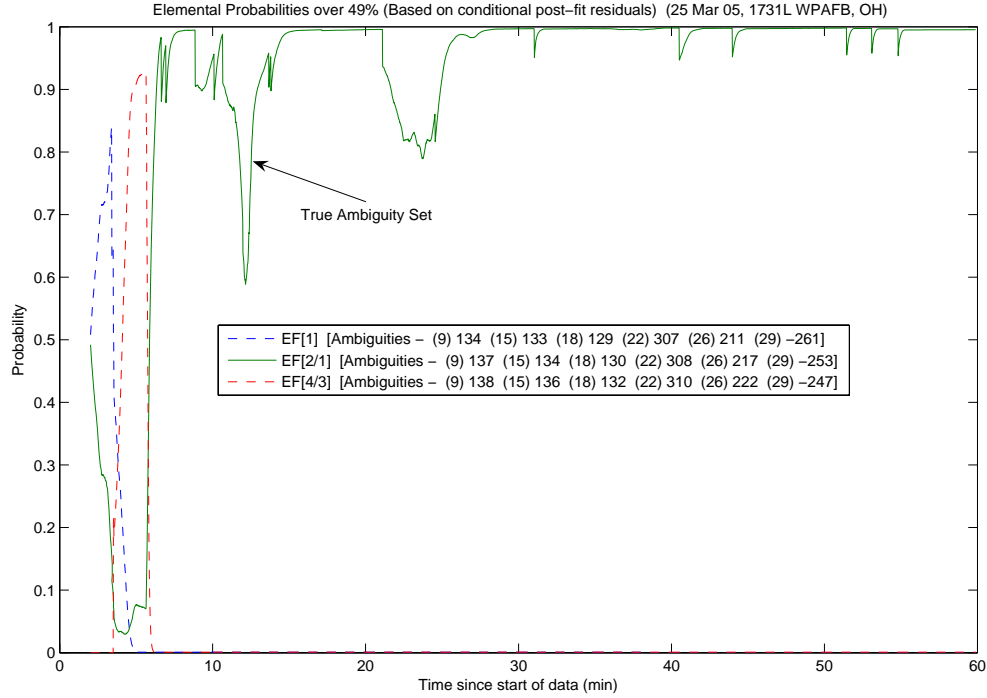


Figure A.14: Case G1: Conditional Probabilities for Selected MMAE Elemental Filters

The SOS residual plot may be difficult to interpret without color. To help identify the true ambiguity set, its SOS residuals are denoted by grey asterisks in all the plots. By 10 minutes into the data run, the true set became apparent to the human eye. As the floating filter continued to gather more information from the measurements, and the covariance matrix “shrunk”, the true set stood out even more, especially after 25 minutes. However, some erroneous sets still managed to approach the true ambiguity set at 42 and 58 minutes in the run. As the satellite geometry changed, certain erroneous ambiguity set’s measurements appeared to “fit” the position solution better due to errors such as multipath and noise. This reduced their SOS residuals to a level equal to, or lower than the true set.

Two important times to note on the residual plot occur at 11 and 22 minutes. At both locations there are erroneous candidate sets whose residuals are lower than the true ambiguity set. In each case, the SOS residuals immediately begin to grow and rise above the true set. These locations correspond to significant dips in the true ambiguity

set's conditional probability as shown in Figure A.14. There are smaller probability dips which correspond to other wrong candidate set's SOS residuals approaching the true set.

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.15 is EF[1], Figure A.16 is EF[2], and Figure A.17 is EF[4].

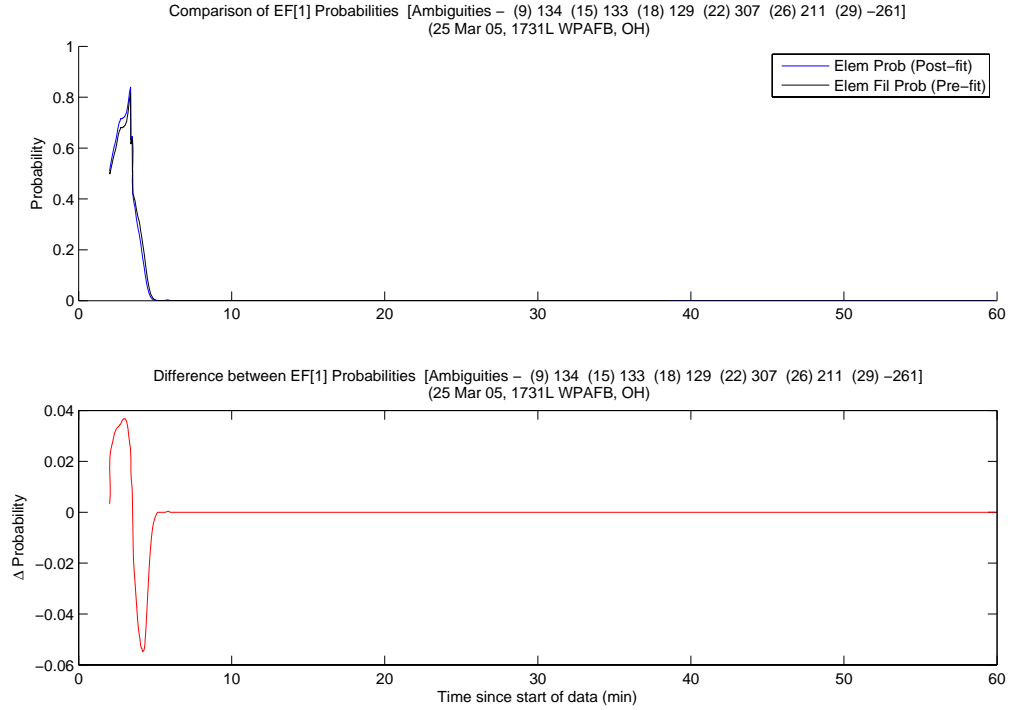


Figure A.15: Case G1: EF[1] Probability Comparison

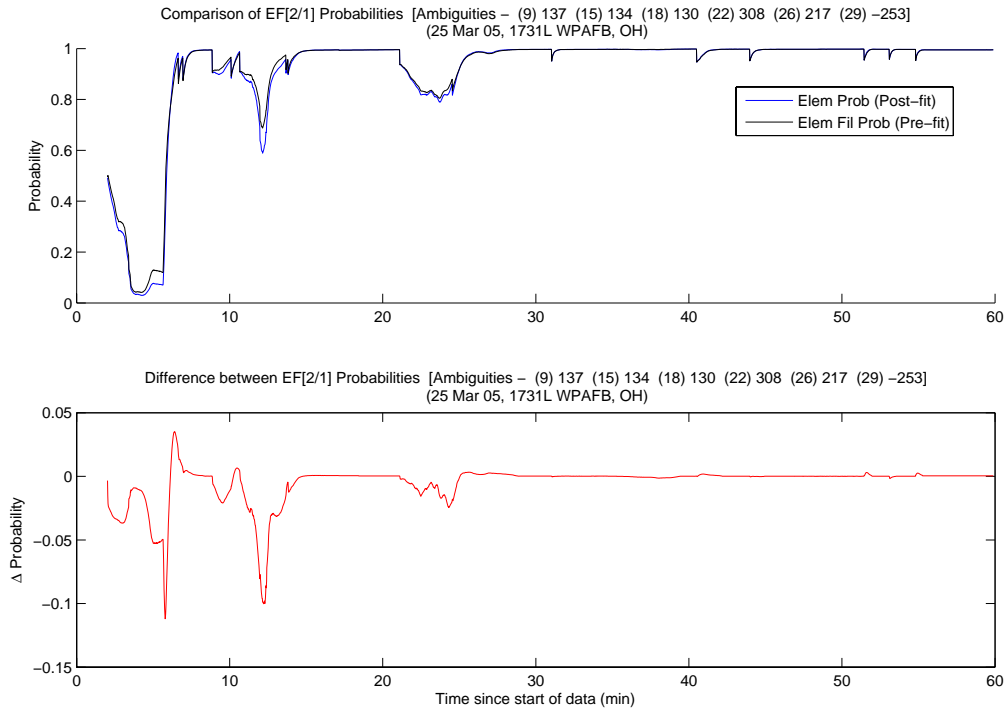


Figure A.16: Case G1: EF[2] Probability Comparison

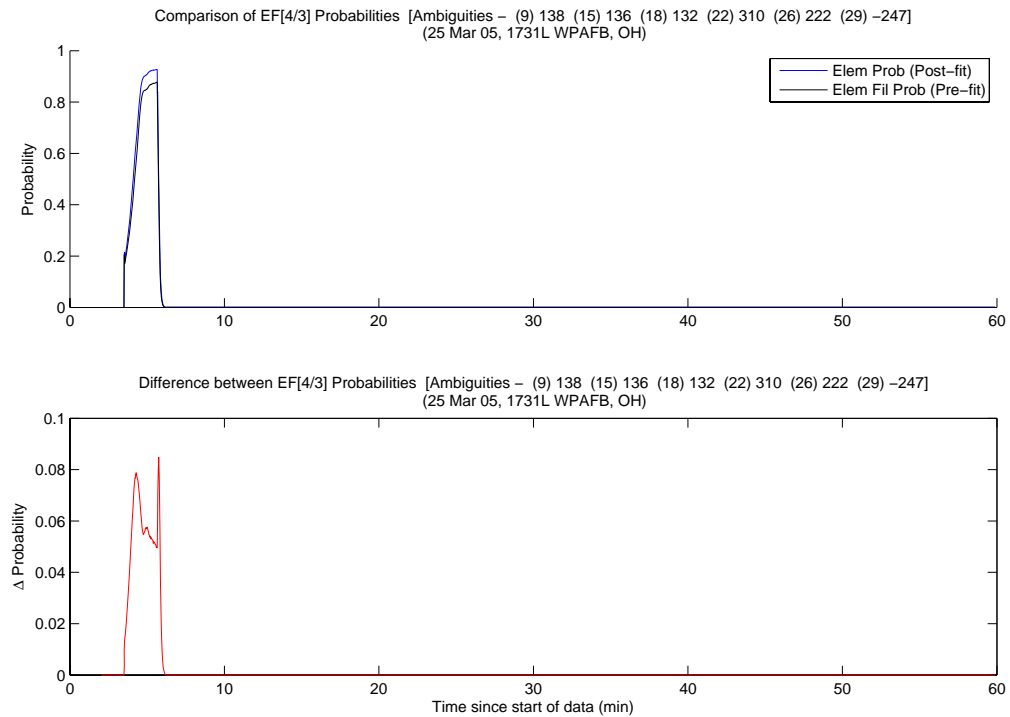


Figure A.17: Case G1: EF[4] Probability Comparison

Dips in the true ambiguity set's probability mean the MMAE blended solution included an incorrect ambiguity set. Since the minimum indicator correctly chose the true set at 4.166 minutes and never lost “lock”, the difference between the MMAE and minimum indicator relative position estimates shows when the blended solution of the MMAE is “pulled” away from the true position. Figure A.18 shows the difference between the MMAE and the minimum indicator position solutions. The figure starts at 4.5 minutes; directly after the true ambiguity set was found by the minimum indicator variable. Any difference between the two positions was due to an incorrect ambiguity set blending with the true ambiguity set. The effect of the reduced probability is visible at 11 and 22 minutes. Also, the minor dips in the probability can be seen in the position error later in the run.

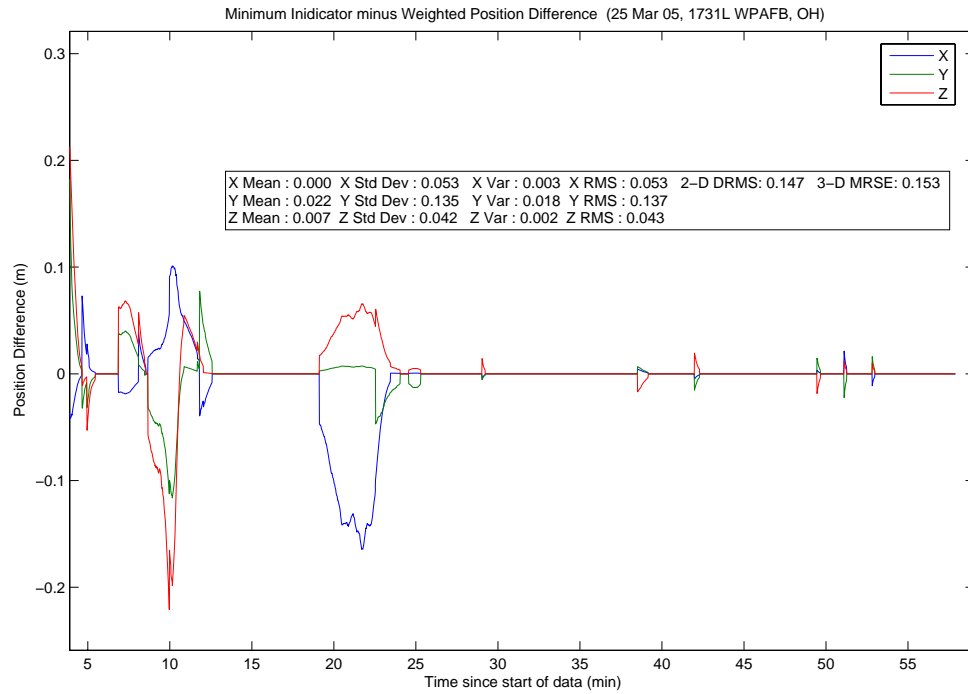


Figure A.18: Case G1: Minimum Indicator and MMAE Position Difference

Finally, Figure A.19 displays the MMAE weighted position error. The MMAE position solution was “pulled” off at 10 and 22 minutes, corresponding to the probability dips in Figure A.14. Figure A.20 shows the relative position error for the minimum indicator variable. The minimum indicator preformed as expected since it was “locked” to the true ambiguity set during the entire data run. The scale on the minimum indicator error plot is one order of magnitude less than the MMAE error plot. Lastly, Table A.1 summarizes the floating point, MMAE and minimum indicator mean position error, standard deviation, 2-D DRMS and 3-D MRSE.

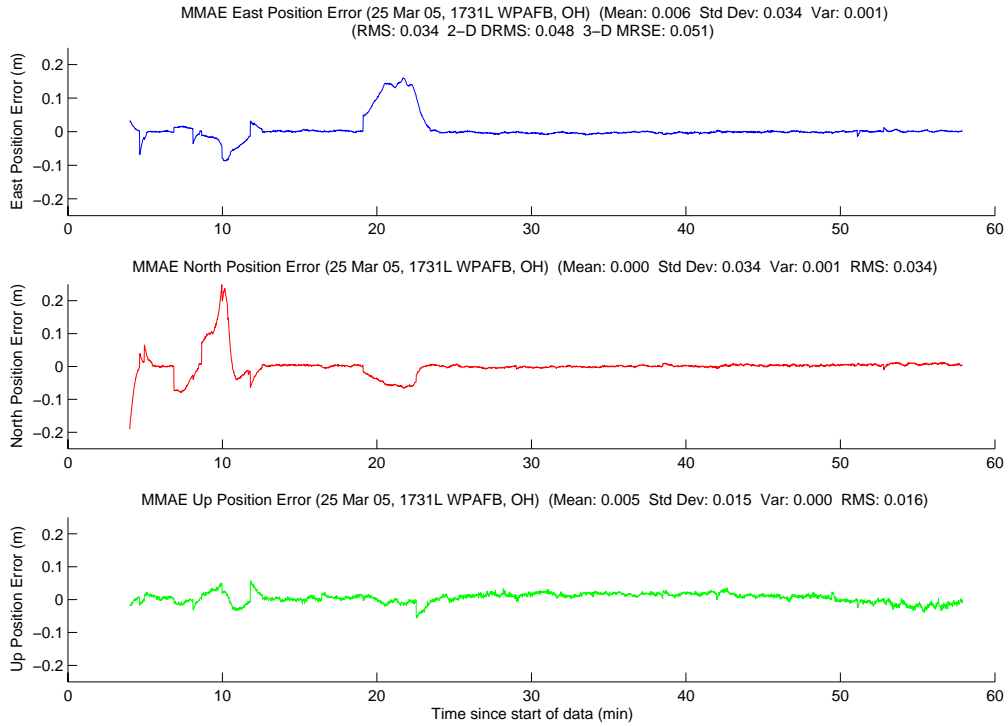


Figure A.19: Case G1: MMAE Position Error

Table A.1: Case G1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.058	0.022	-0.023	0.036	-0.036	0.046	0.075	0.095
MMAE	0.006	0.034	0.000	0.034	0.005	0.015	0.048	0.051
Min. Ind.	0.001	0.003	0.002	0.003	0.007	0.011	0.005	0.014

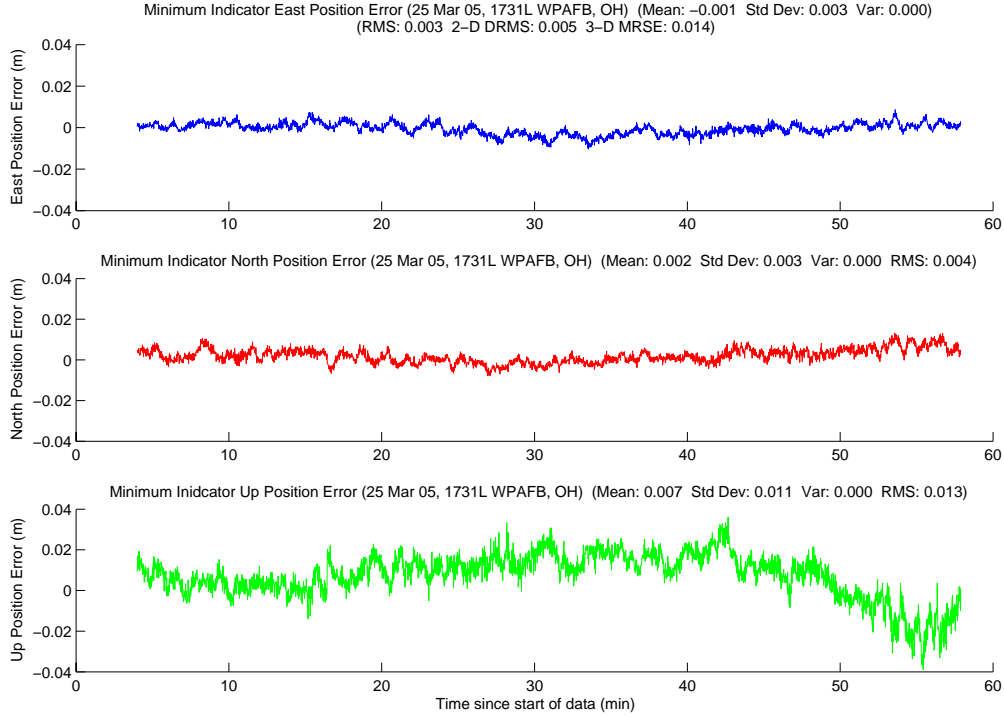


Figure A.20: Case G1: Minimum Indicator Position Error

A.1.2 Case G2: Dynamic Driving Test. The second ground test was a dynamic driving test conducted on 30 March 2005. During the test, the cars ranged from 3.5 to 77.6 meters of separation. This test shows one of the single-cycle slips mentioned in Chapter III, and its effect on the floating filter, MMAE and minimum indicator. Once again, the minimum indicator selected the true ambiguity set immediately at four minutes. There were no other sets chosen before the next true ambiguity set was discovered at 50.8 minutes into the data run. The true ambiguity sets were the only sets selected by the minimum indicator variable.

Throughout the data run, there were a total of 39 unique ambiguity sets returned from the LAMBDA algorithm. The first true ambiguity set was the top candidate set returned from the LAMBDA algorithm on the first call. Later in the run, LAMBDA returned the second true ambiguity set 30 seconds after the cycle slip was detected. In all, the true ambiguity set was returned as the top set from LAMBDA, 2984 of 3474 epochs (85.9 percent), and reached a maximum ratio of 4.4.

The floating filter performance is shown in Figures A.21, A.22, and A.23. The floating filter had trouble converging initially. The most likely cause was the lower number of visible satellites when the filter started. The filter initialized with seven satellites, but then lost a satellite three minutes later. Six satellites were about the minimum required for the floating filter to operate efficiently with only L1 observables. At 18 minutes into the data run, the filter began converging properly. This coincides with two satellites rising into view within three minutes of each other; denoting the importance of the number of visible satellites and their associated geometry.

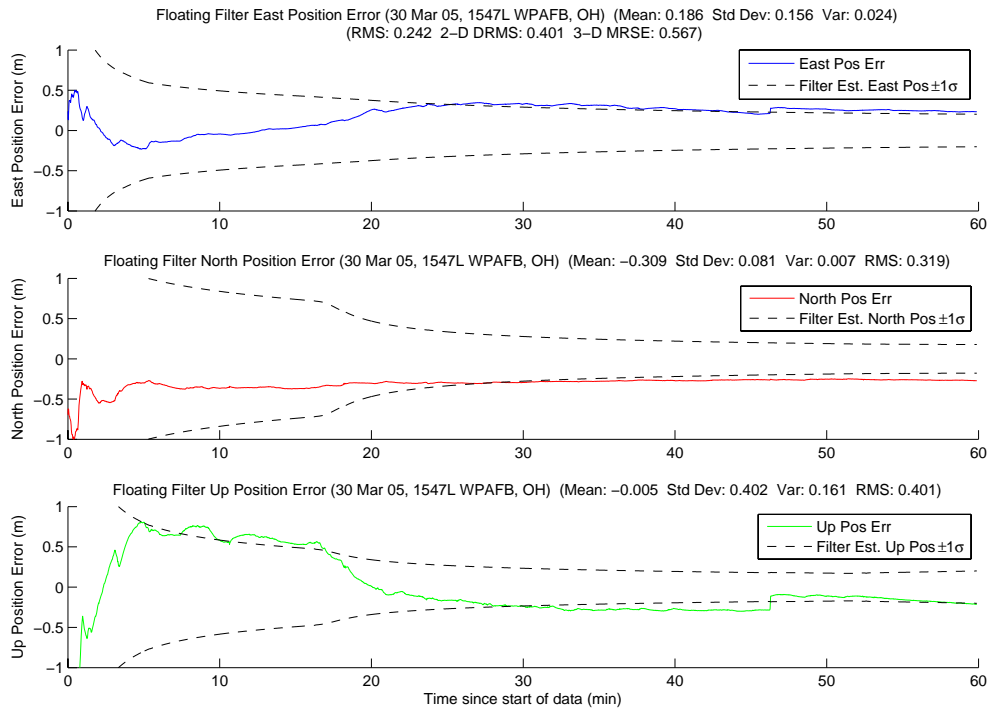


Figure A.21: Case G2: Floating Filter Position Errors

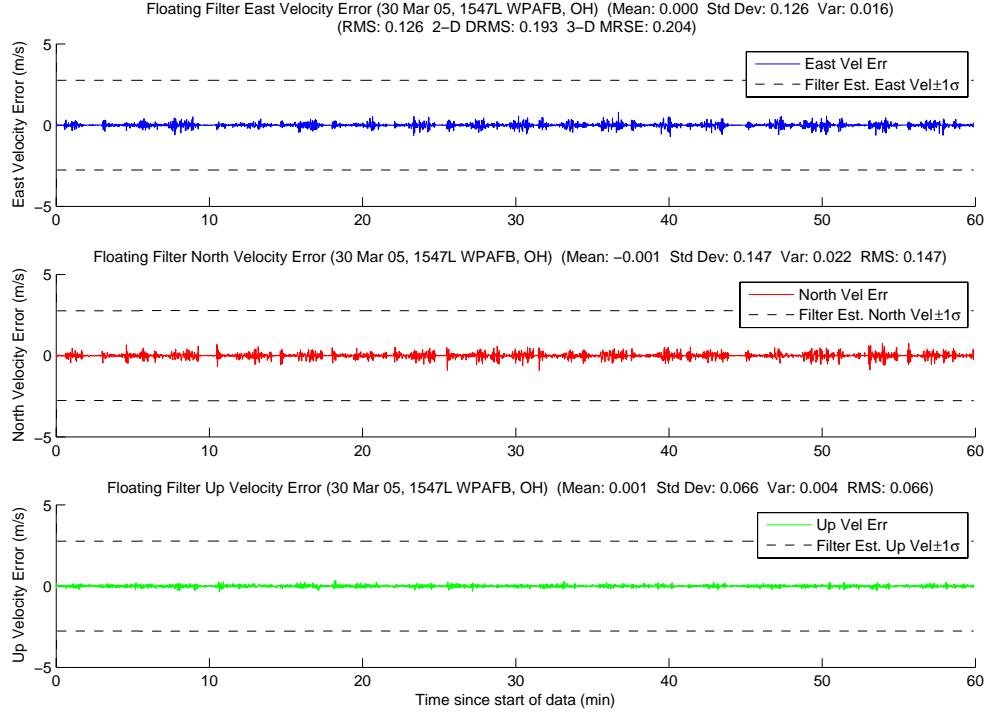


Figure A.22: Case G2: Floating Filter Velocity Errors

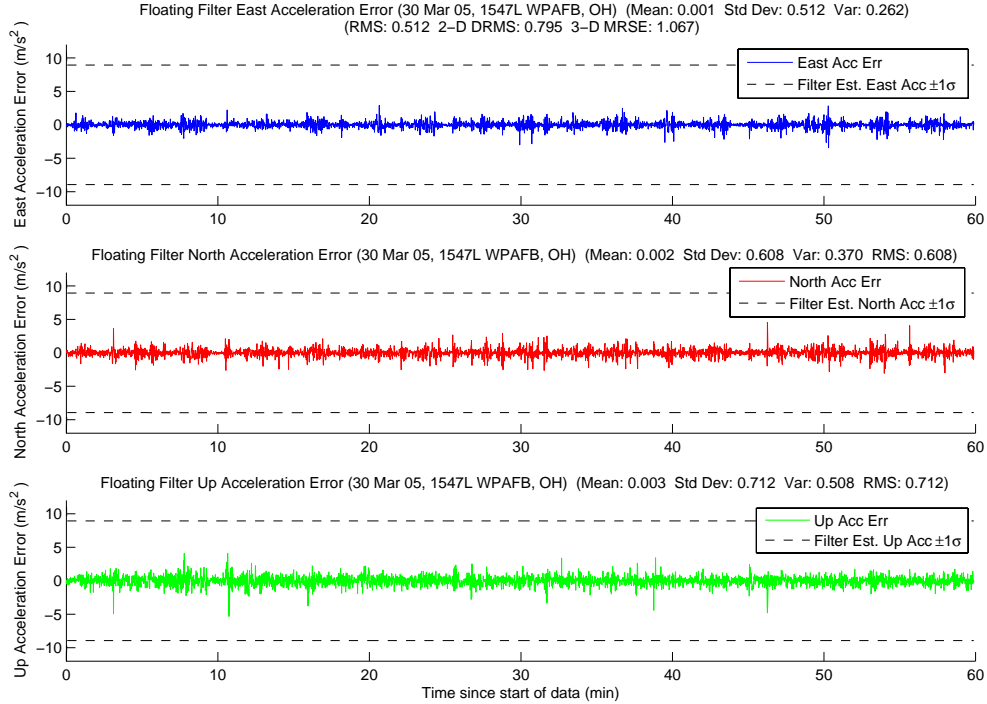


Figure A.23: Case G2: Floating Filter Acceleration Errors

Figure A.24 displays PRN 6 measurement information and the floating filter's ambiguity estimate. The cycle slip that caused the minimum indicator problems is visible at approximately 46 minutes with a “jump” in the floating filter ambiguity error. The floating filter's 1σ estimate of the ambiguity state converged in one epoch. Thus, there was only a tiny bump in the $\mathbf{P}_k(t_i^+)$ 1σ value at the cycle slip, even though PRN 6 variance was reset. Finally, the phase residuals were significantly higher than those from Case G1, in Figure A.4 (note scale difference). The higher phase residuals were due to the vehicle dynamics. The rest of the satellite's plots follow in Figures A.25, A.26, A.27, A.28, A.29, A.30, A.31, and A.32.

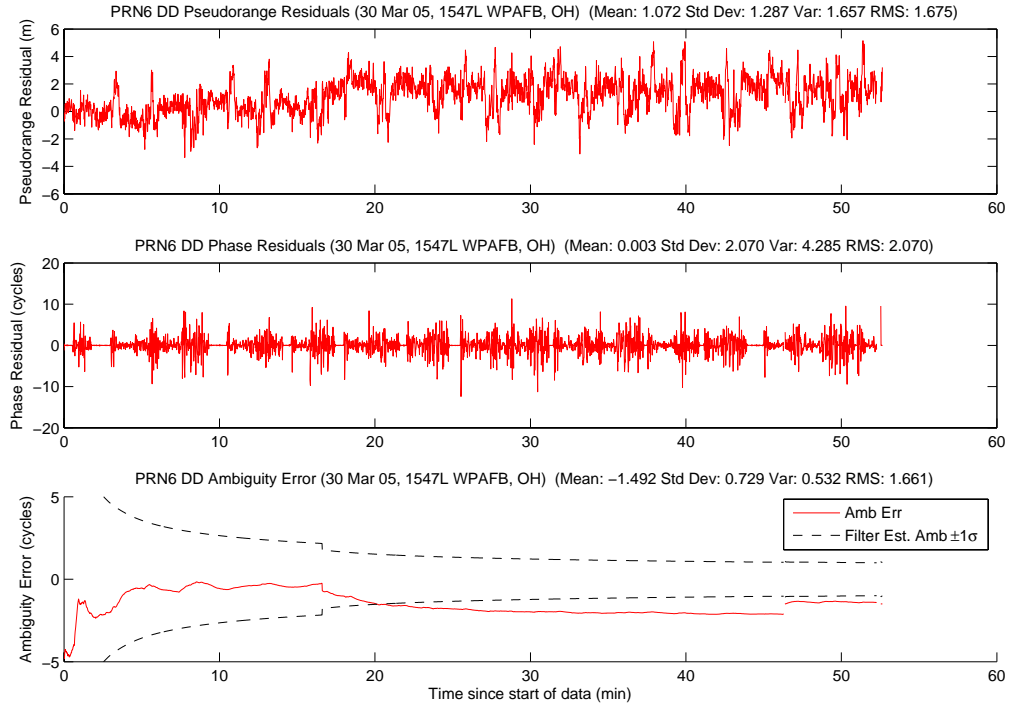


Figure A.24: Case G2: Satellite 6 Measurements

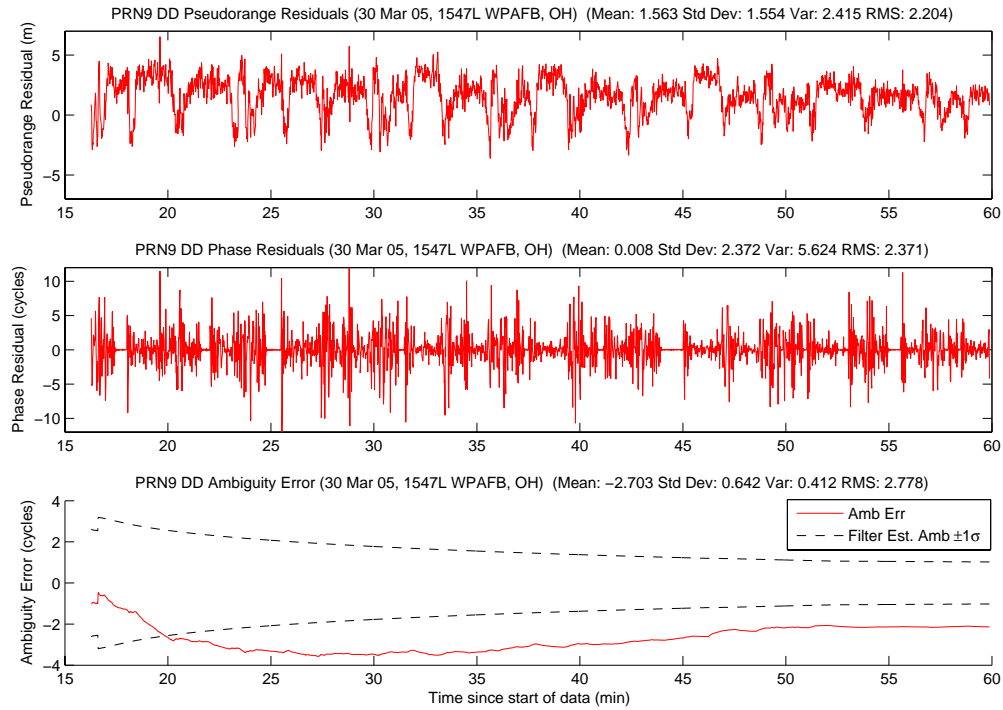


Figure A.25: Case G2: Satellite 9 Measurements

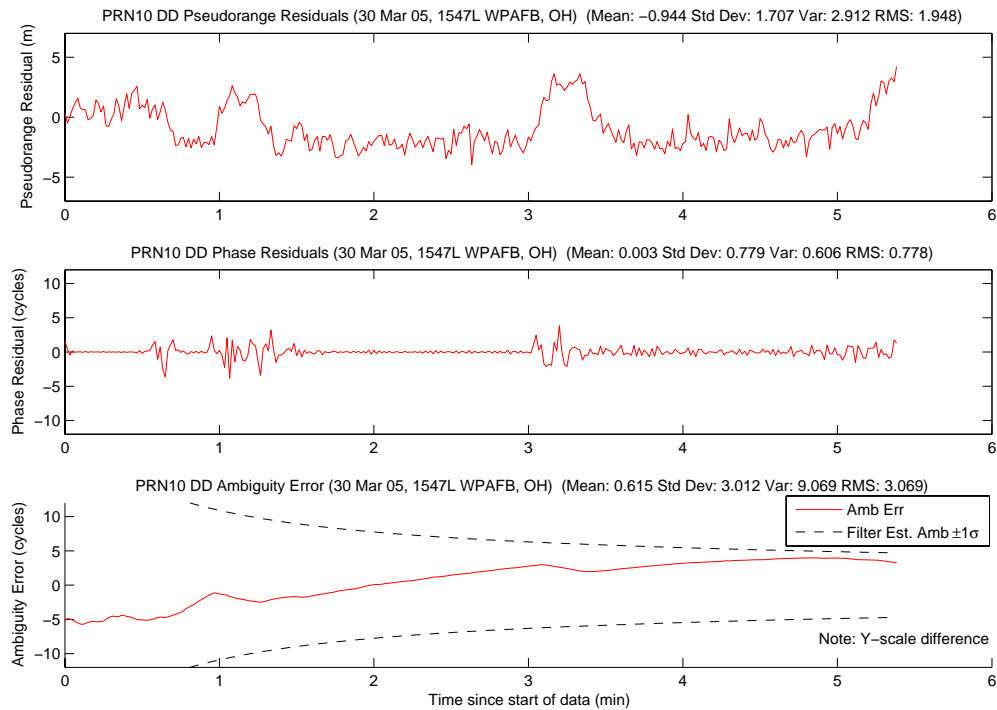


Figure A.26: Case G2: Satellite 10 Measurements

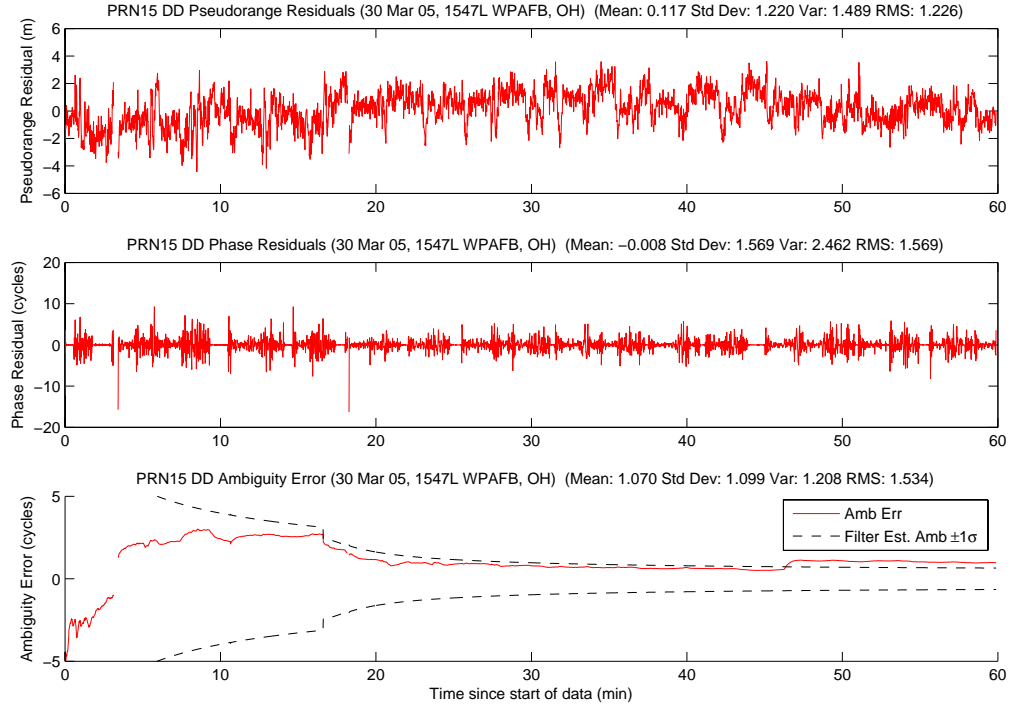


Figure A.27: Case G2: Satellite 15 Measurements

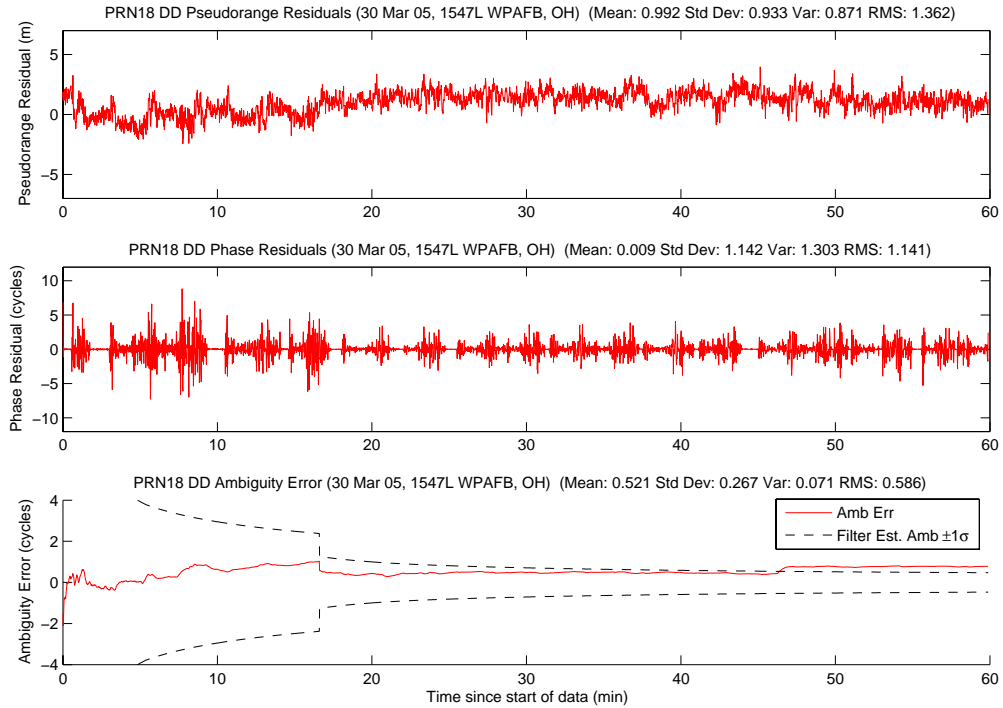


Figure A.28: Case G2: Satellite 18 Measurements

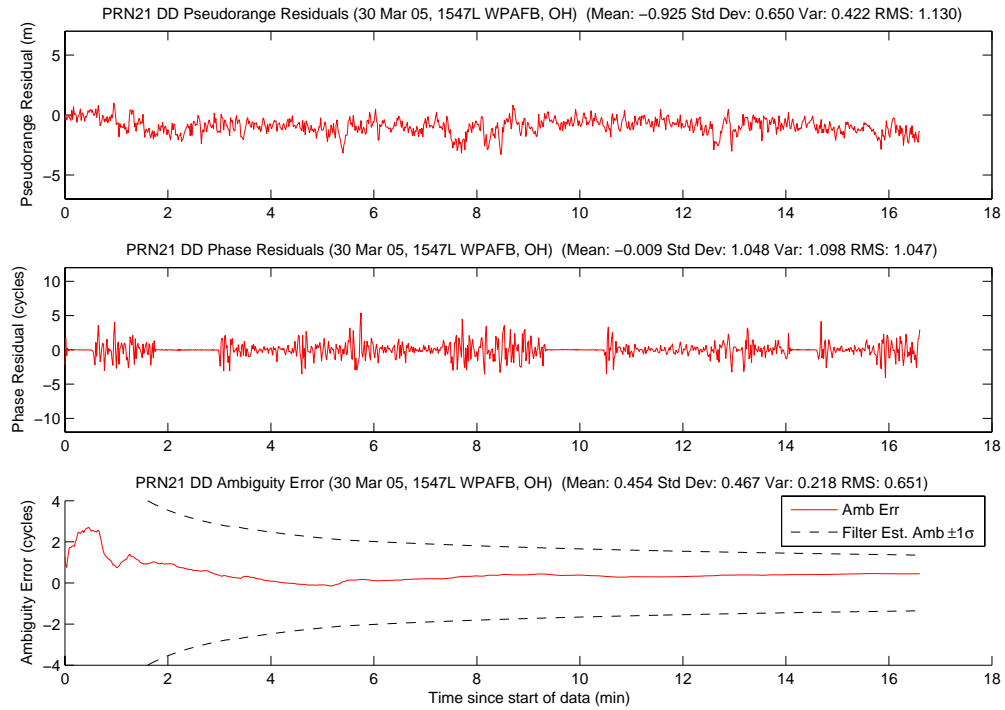


Figure A.29: Case G2: Satellite 21 Measurements

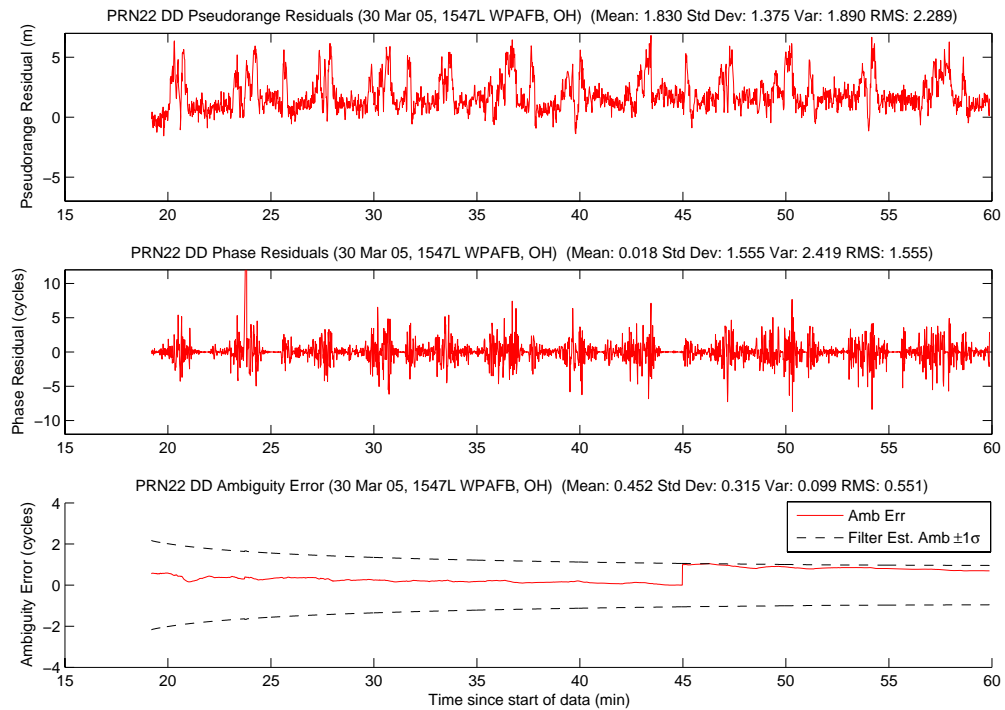


Figure A.30: Case G2: Satellite 22 Measurements

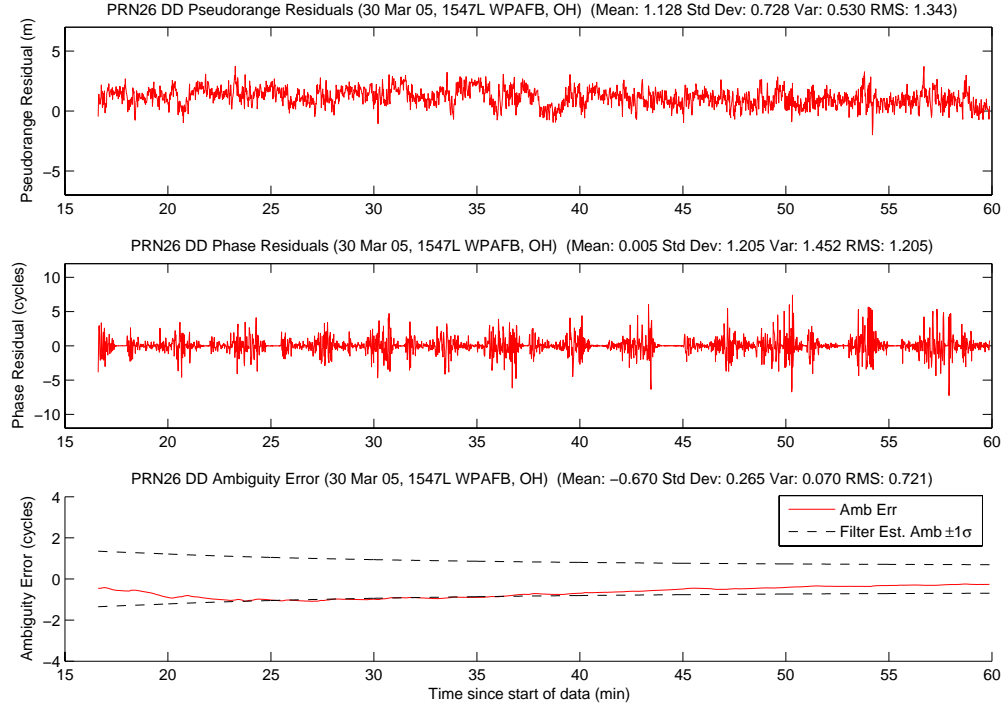


Figure A.31: Case G2: Satellite 26 Measurements

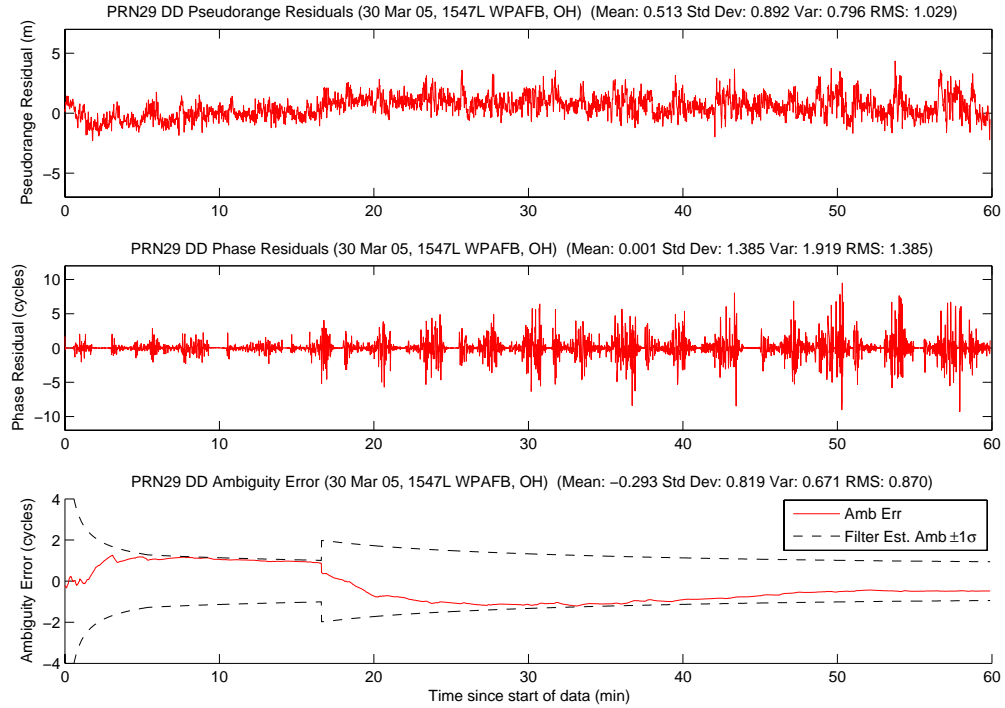


Figure A.32: Case G2: Satellite 29 Measurements

Figure A.33 shows the number of visible satellites. The satellite dropouts were due to cycle slips (real or erroneous) or loss-of-lock. The only satellite vehicle (SV) dropouts due to actual cycle slips were at approximately 46 and 51 minutes. PRN 6 had a cycle slip at 46 minutes, which initially caused the minimum indicator to be incorrect. A few minutes later, PRN 22 also had a cycle slip, but was reset before the minimum indicator determined the true ambiguity set.

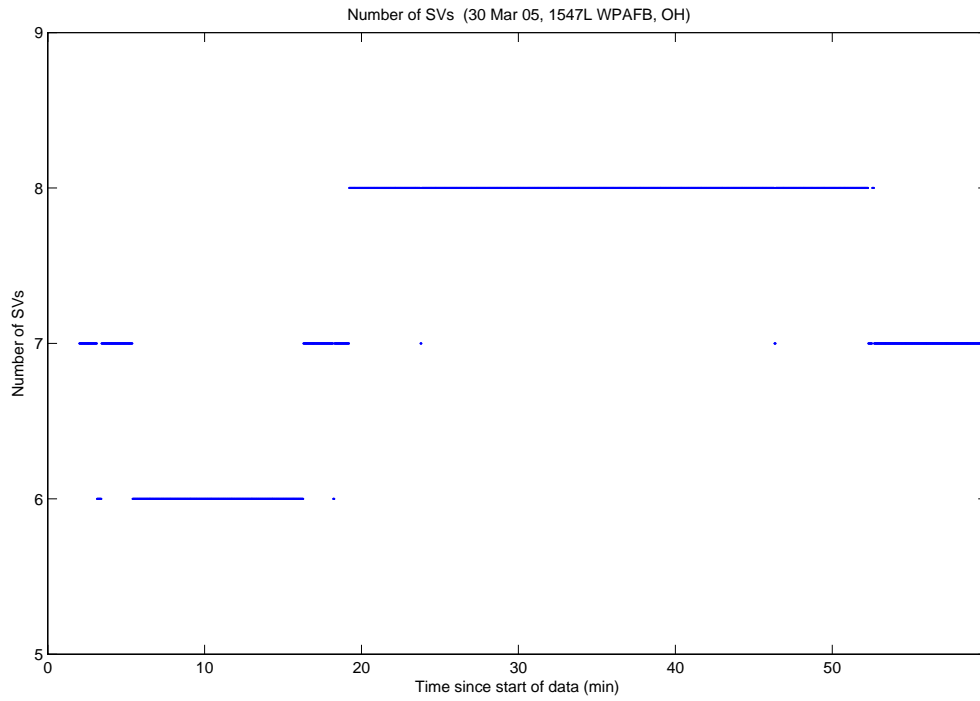


Figure A.33: Case G2: Visible Satellites

The residual plot, Figure A.34, is expanded to better show the true ambiguity set's residuals. The probability plot, Figure A.35, has three points of interest. The first is at approximately 3 minutes where the MMAE determined the true ambiguity set. This was a full minute before the minimum indicator. Then at 10 minutes, another false candidate set absorbed 50 percent of the conditional probability, which corresponds to a rise in the true set's residuals in Figure A.34. Finally, at 46.33 minutes the true ambiguity set's probability immediately dropped when the cycle slip occurred. After the cycle slip, numerous erroneous ambiguity sets vied for the released probability. After 30 seconds, the second true ambiguity set was returned from LAMBDA and absorbed all of the probability a minute later as seen in Figure A.36.

The minimum indicator did not select the second correct set until 50.8 minutes into the run, about four minutes after the cycle slip. The 30 seconds for true ambiguity set to be returned from LAMBDA, coupled with the initial 2 minute SOS residuals acquisition time, means it took the minimum indicator 1.5 minutes to find the true set. The 1.5 minute delay was primarily driven by the fact that the minimum indicator divides by the summation of total number of times a candidate was returned as the top set from LAMBDA and the number of epochs it has been "alive". Since the first true ambiguity set had been the top set over 2200 epochs and was "alive" for approximately 3200 epochs, it took time for its minimum indicator value to grow (based on its much higher residuals) above the new true ambiguity set. Time constraints limited the author's ability to address this problem in the code, but a potential fix is discussed in Chapter V.

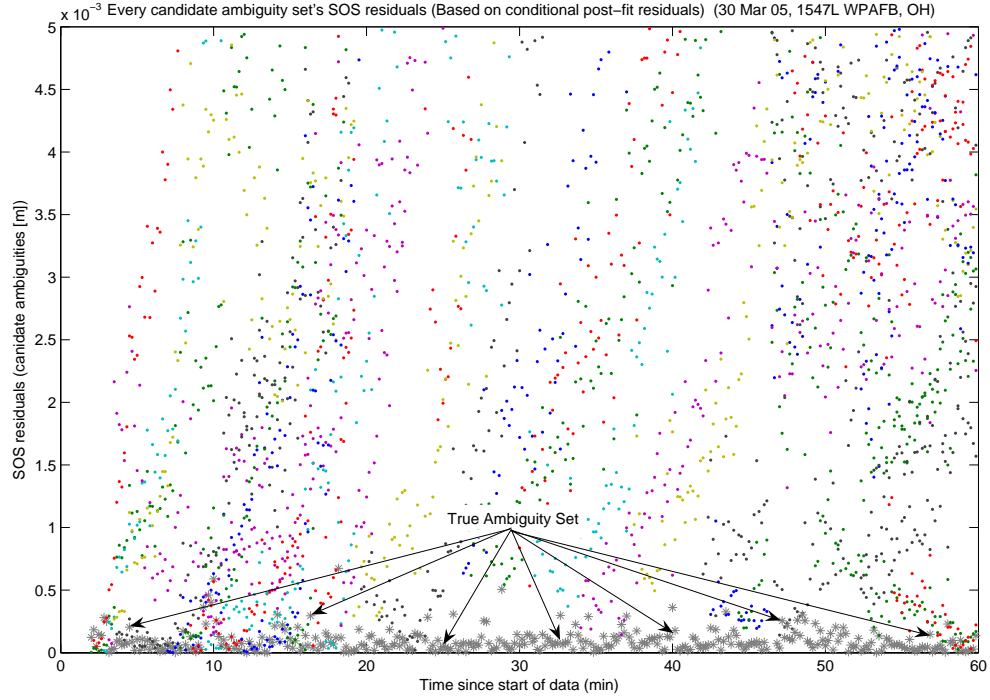


Figure A.34: Case G2: SOS Residuals

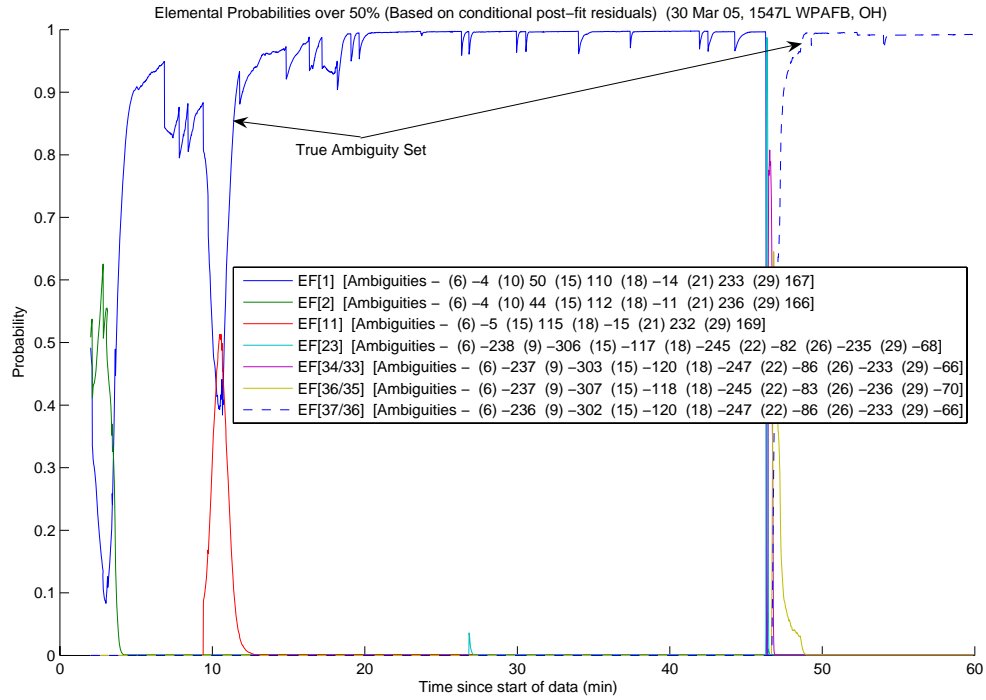


Figure A.35: Case G2: Conditional Probabilities for Selected MMAE Elemental Filters

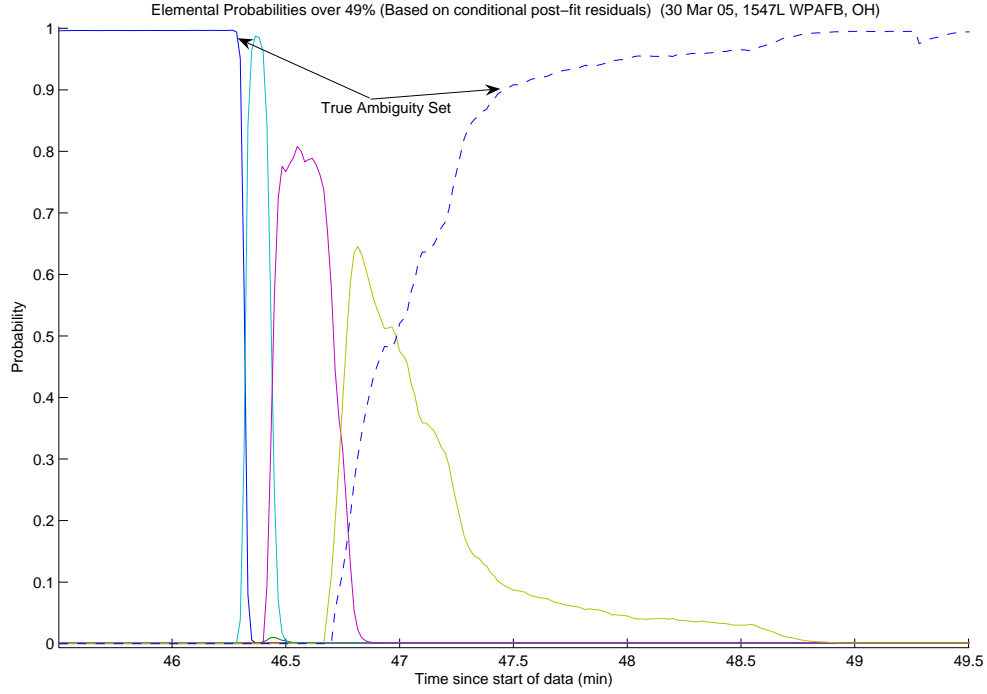


Figure A.36: Case G2: Conditional Probabilities for Selected MMAE Elemental Filters Expanded

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.37 is EF[1], Figure A.38 is EF[2], Figure A.39 is EF[23], Figure A.40 is EF[34], Figure A.41 is EF[36], and Figure A.42 is EF[37].

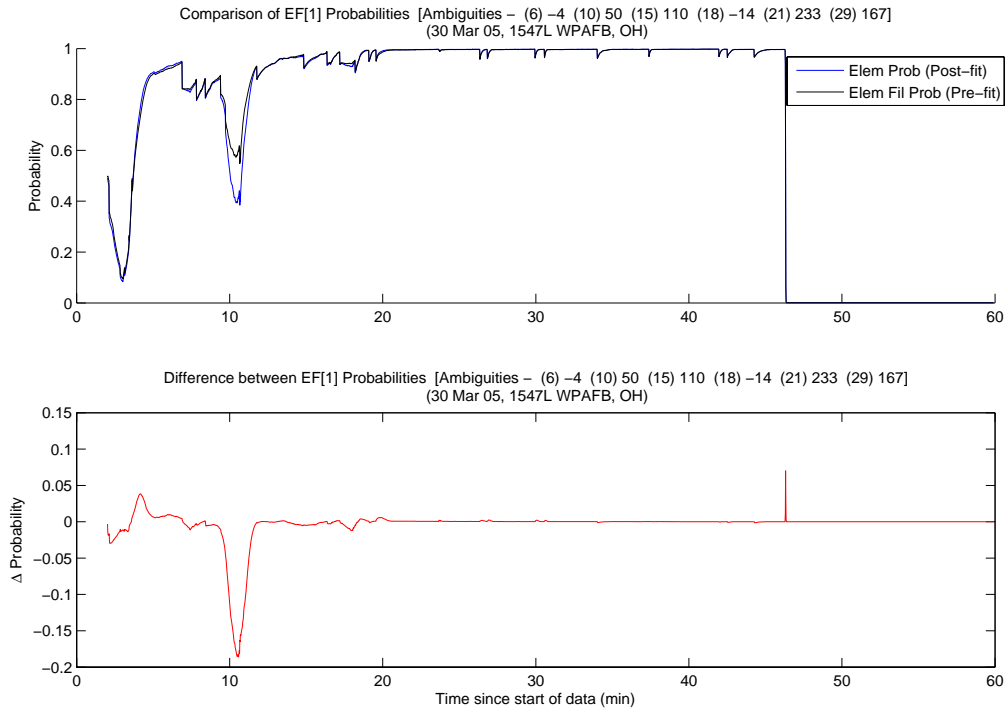


Figure A.37: Case G2: EF[1] Probability Comparison

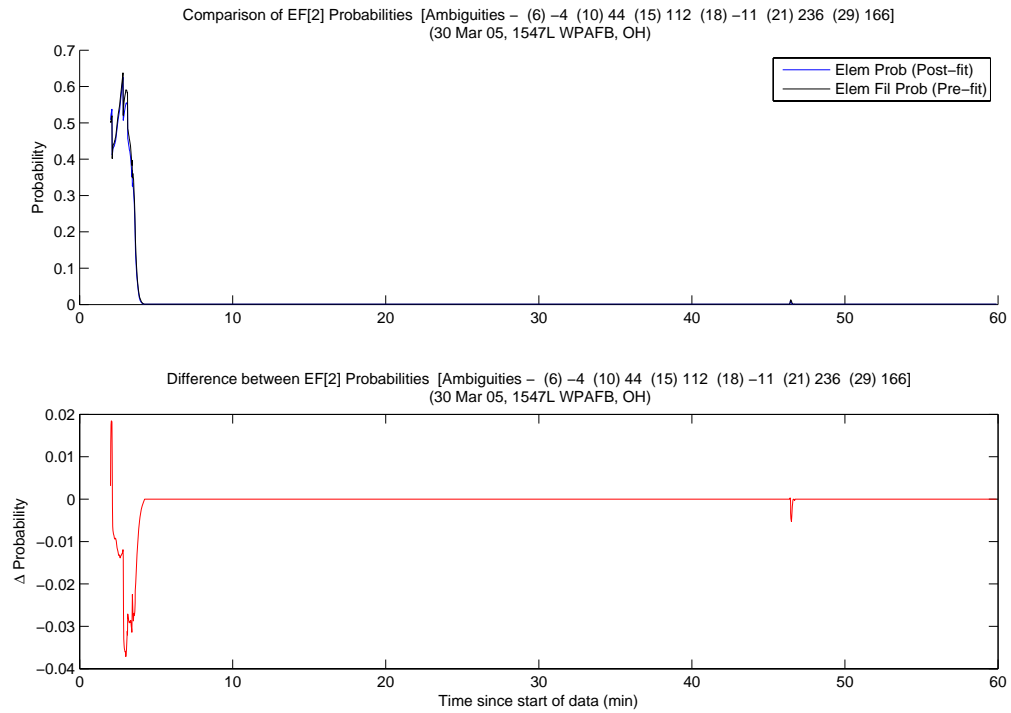


Figure A.38: Case G2: EF[2] Probability Comparison

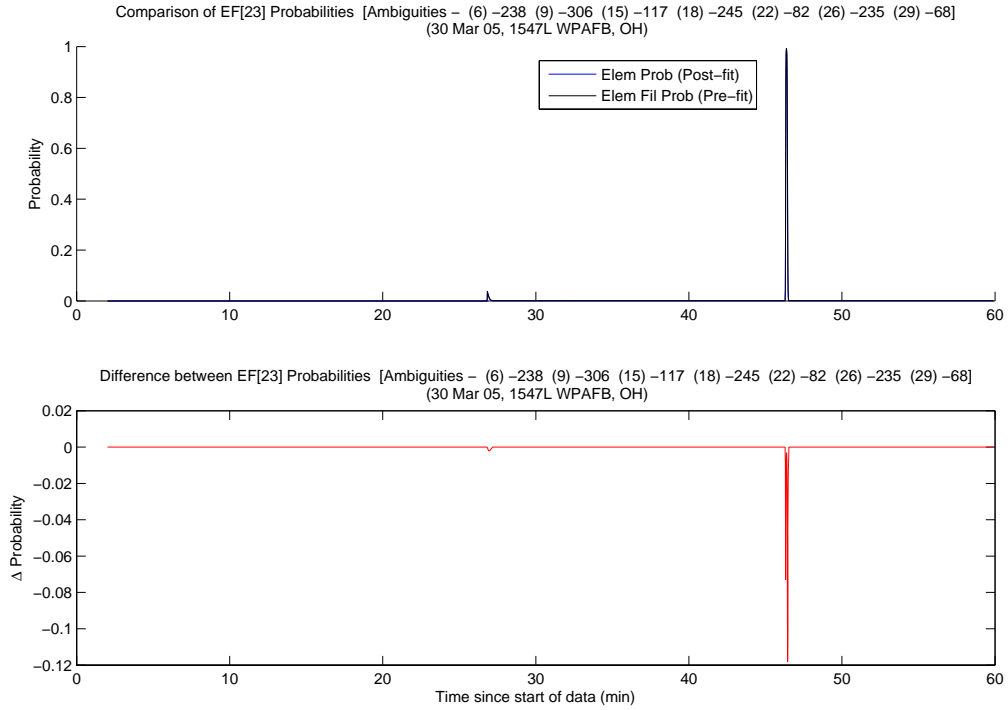


Figure A.39: Case G2: EF[23] Probability Comparison

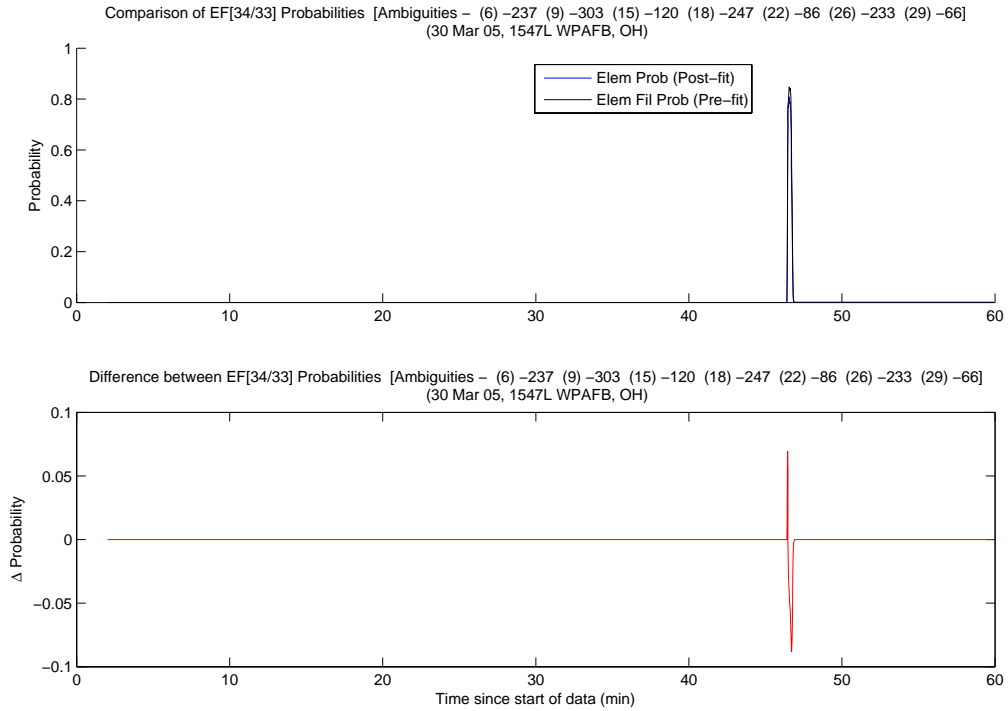


Figure A.40: Case G2: EF[34] Probability Comparison

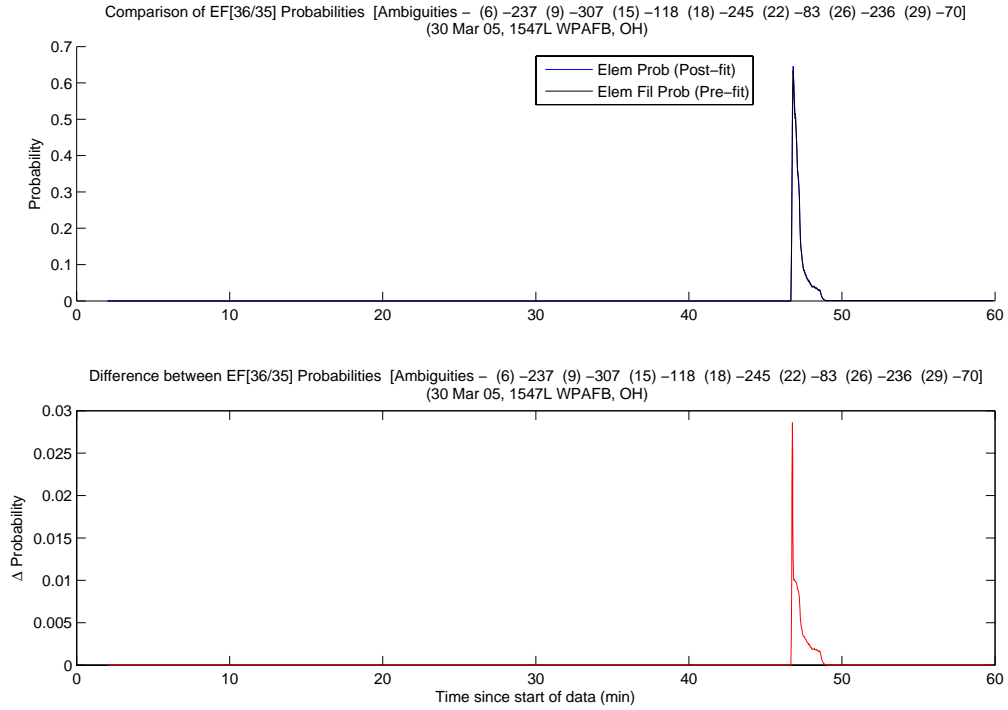


Figure A.41: Case G2: EF[36] Probability Comparison

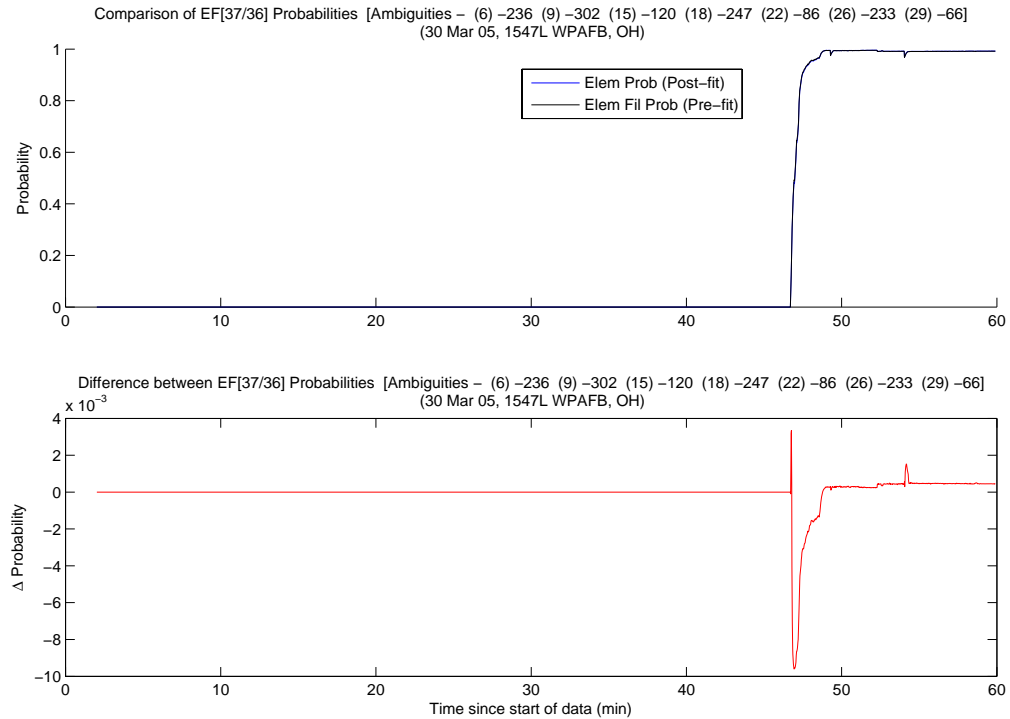


Figure A.42: Case G2: EF[37] Probability Comparison

Figure A.43 shows the MMAE position error. Since the MMAE quickly determined the true ambiguity set, it is relatively accurate from the beginning of the data run unlike the other ground case. However, the probability dip at 10 minutes and the cycle slip at 46 minutes can both be seen in the plot.

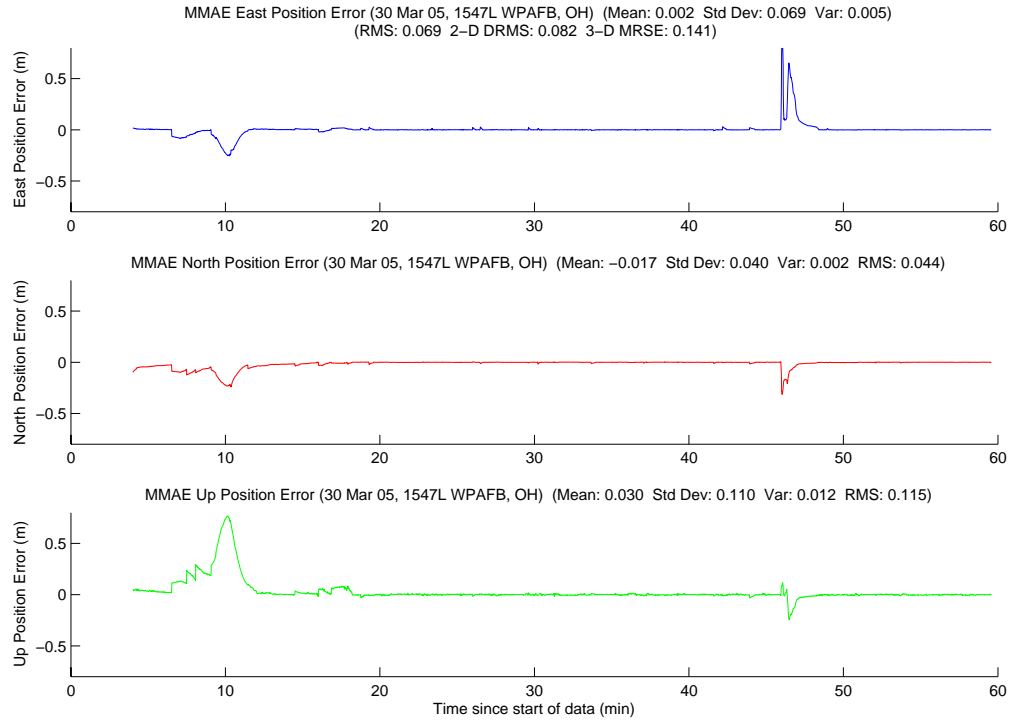


Figure A.43: Case G2: MMAE Position Error

Finally, Figure A.44 displays the position error plot for the minimum indicator variable and Table A.2 gives a summary of each method's position error. Even with the cycle slip, the minimum indicator remained reasonably accurate. Since the filter was using L1 observables, and only one satellite had a single-cycle cycle slip, the relative position error remained small. The magnitude of the error was approximately 18.7 cm. As expected, the error was close to the L1 wavelength of 19.0 cm. The reduced error from cycle slips and slightly better overall accuracy are the main advantages of L1 observables over widelane observables.

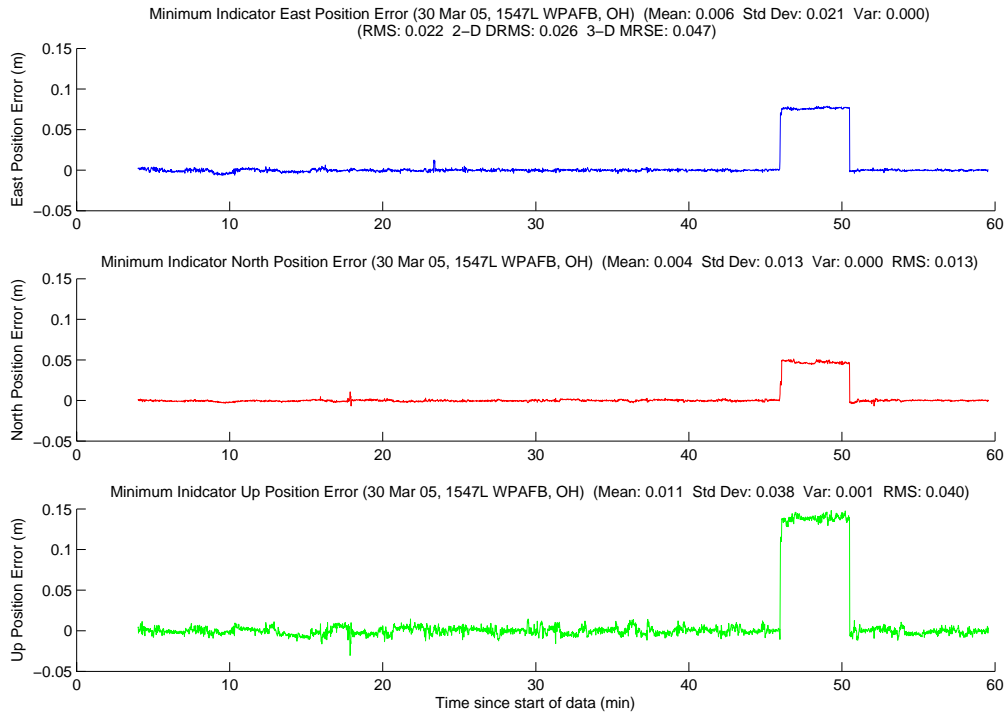


Figure A.44: Case G2: Minimum Indicator Position Error

Table A.2: Case G2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.186	0.156	-0.309	0.081	-0.005	0.402	0.401	0.567
MMAE	0.002	0.069	-0.017	0.040	0.030	0.110	0.082	0.141
Min. Ind.	0.006	0.021	0.004	0.013	0.011	0.038	0.026	0.047

A.1.3 Ground Summary. Two ground tests were discussed with a total run time of approximately two hours. During the two hours, there were five SV dropouts, all during the dynamic driving test. Two were due to actual cycle slips and three were due to erroneously detected cycle slips. The algorithm correctly handled all of the SV dropouts. The base satellite changed twice during the two runs. Once more, the algorithm correctly handled the base satellite changes. Most importantly, all ground objectives were met. The system operated in realtime, the datalink was verified to useful ranges, the algorithm successfully resolved the true ambiguities, and appropriate tuning values were determined for the Kalman filter. The system was declared ready for flight test.

A.2 Flight Tests

Seventeen flight test data runs were performed for a total of approximately 14 hours. The algorithm did not crash a single time during the 14 hours of flight test. Even when messages were “lost”, satellites dropped out, the base satellite changed, or cycle slips occurred, the algorithm handled all the contingencies. There were a total of 91 satellite dropouts during actual flight testing. No dropouts while the aircraft were on the ground were included, since all of the dropouts on the ground were due to missed messages (the antennas did not have clear line-of-sight). Of the 91 dropouts, 2 were due to erroneously detected cycle slips, 6 were actual cycle slips, and the remaining 83 dropouts were due to masking. Also, the base satellite changed 10 times over the 14 hours of flight test. Table A.3 summarizes the satellite information. Except when long strings of messages were “lost”, causing erroneous cycle slips on all of the satellites, the algorithm correctly handled every contingency.

Not only did the algorithm properly handle contingencies, it also had good performance. The minimum indicator met the centimeter-level requirement 13 of the 17 data runs, four of which were in the millimeter range. The times when the minimum indicator did not meet the centimeter-level requirement, it incorrectly chose an erroneous ambiguity set only for a portion of the run. Table A.4 gives a summary

Table A.3: Satellite Summary for Flight Testing

Total Data Sets	17 sets
Total Hours	14 hours
Base Satellite Changes	10 changes
Total Satellite Dropouts	91 dropouts
Erroneous Cycle Slips	2
Actual Cycle Slips	6
Satellites Masks	83

of the DRMS and MRSE for the MMAE and minimum indicator on each run. Runs in bold are times when the minimum indicator did not meet the centimeter requirement. Case F4.1 was the only run the that the MMAE performed better than the minimum indicator. The reason is discussed later in Section A.2.7. Also, an average is shown in the last row.

Table A.4: Flight Test Position Error Summary (m)

Data Set	Min. Ind.		MMAE	
	DRMS	MRSE	DRMS	MRSE
F1.1	0.274	0.775	0.837	1.333
F1.2	0.041	0.085	0.363	0.417
F2.1	0.044	0.062	0.116	0.121
F2.2	0.004	0.023	0.343	0.738
F3.1	0.135	0.137	0.205	0.246
F3.2	0.090	0.153	0.192	0.229
F4.1	0.179	0.243	0.055	0.068
F5.1	0.002	0.011	0.004	0.013
F5.2	0.006	0.050	0.059	0.131
F6.1	0.003	0.010	0.074	0.138
F6.2	0.002	0.010	0.021	0.023
F6.3	0.004	0.007	0.012	0.021
F6.4	0.002	0.007	0.021	0.027
F7.1	0.002	0.008	0.011	0.015
F7.2	0.002	0.007	0.030	0.037
F8.1	0.003	0.012	0.054	0.072
F8.2	0.006	0.016	0.017	0.023
Average	0.047	0.095	0.142	0.215

Table A.4 is a good summary of the errors during each data run. However, it does not tell the entire story. The DGPS system designed in this thesis was for use when the aircraft were in close formation (within 500 feet). Although interesting from an academic stand point, the times when the aircraft were outside this range should not be included in the error. For this reason, Figure A.45 was constructed which displays every epoch when the aircraft were within 500 feet of each other. Each data run is separated by black lines with the appropriate date and time annotated on the plot. The x-axis shows the time from each data run (i.e., a time referenced within a particular data run in Figure A.45 matches the time in the minimum indicator plots displayed later). The 1103L data run on 27 April was the only set where the aircraft were not within 500 feet.

Over the course of approximately 11 hours, the DRMS and MRSE for the combined data runs was 1.9 and 3.3 centimeters, respectively! There was only one 40 second time period when the minimum indicator was not using the true ambiguity set. This was due to intermittent datalink problems which caused a system reset during the 1020L data run on 05 October. Excluding this 40 second period from the data, results in a remarkable DRMS of 0.3 centimeters and MRSE of 1.0 centimeter. The system performed almost flawlessly in its designed operational envelope.

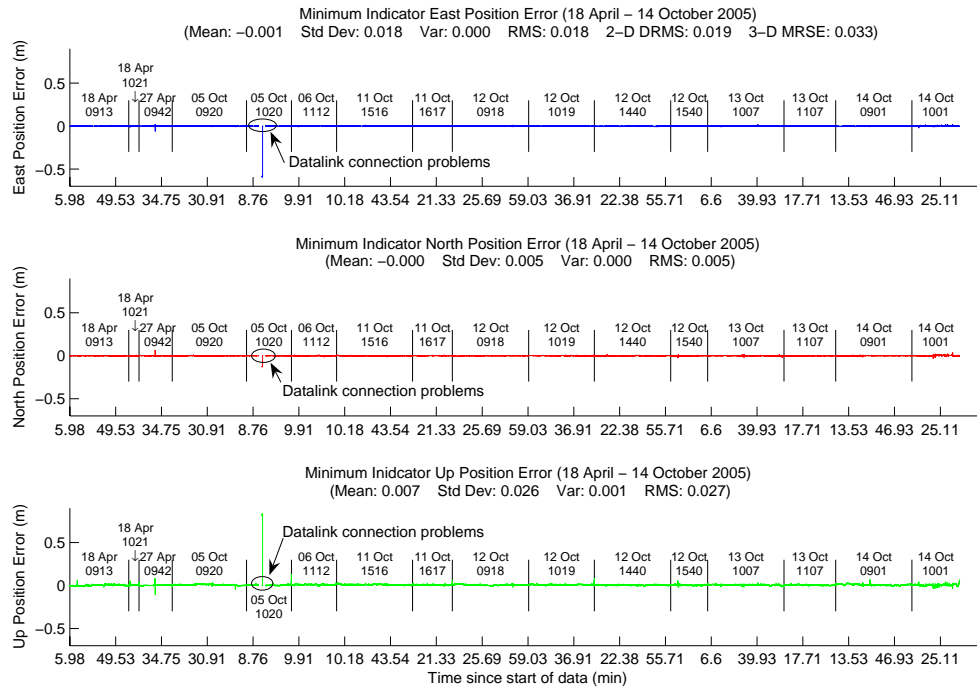


Figure A.45: Every Epoch When the Aircraft Were Within 500 Feet

The first four cases were from the April flight tests and the remaining sets were from the October flight tests. The first set of tests in April were between two C-12Cs. These tests verified the data-link capability and DGPS algorithm position solution. The second set of tests in October were between a C-12C and Learjet 24. The Learjet had a modified flight control system that allowed the test team to input its own flight control laws for simulated aerial-refueling. The non-inertial hardware was identical between the April and October tests. Each flight test was divided into two data runs of approximately one hour each. Due to hardware problems, cases F2.1 and F2.2 were only 43 and 23 minutes long, respectively.

A.2.1 Case F1.1: Flight Test 1, First Data Set. Case F1.1 was flown at EAFB on 18 April 2005. The flight took off at 0924L, and the weather was clear with some gusty winds. The actual data set begins just after engine start, runs through take-off and into the area work. During the test, the aircraft ranged from 18.8 to 1448.2 meters apart. Figure A.46 shows the minimum indicator North, East, Down relative positions. The data run went well except for take-off.

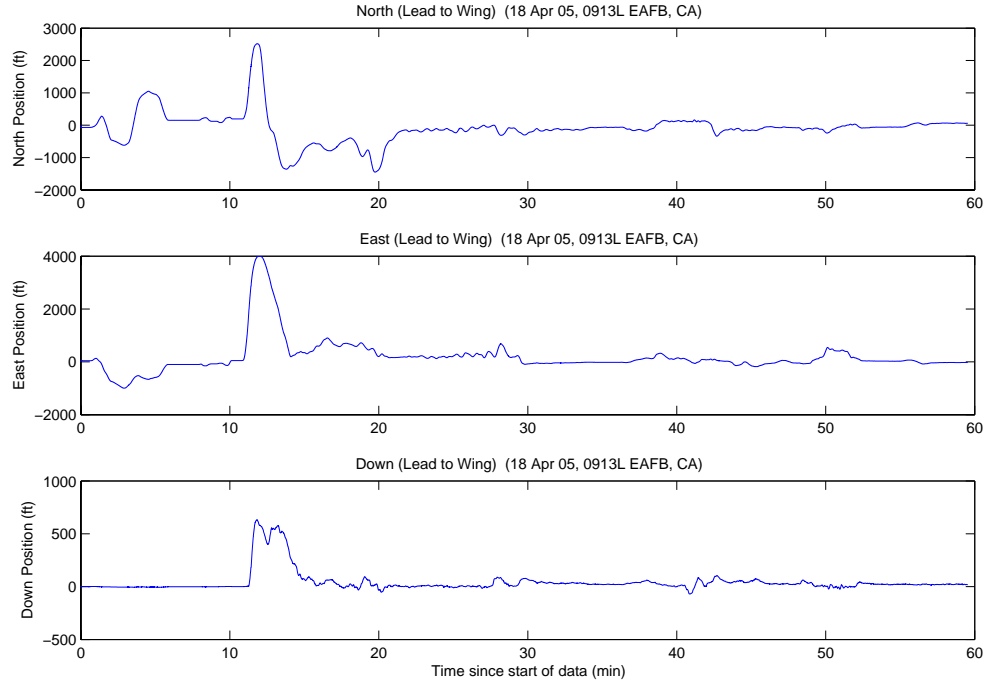


Figure A.46: Case F1.1: North, East, Down Relative Positions

For an unknown reason the floating filter reset approximately 50 feet above the ground during take-off. The lead aircraft had just begun a turn, masking the data-link antenna from the wing aircraft. During post-processing the floating filter did not reset, however five messages from the lead aircraft were lost at the exact time of the in-flight filter reset. Most likely, the algorithm incorrectly detected cycle slips on all of the satellites and set a reset flag.

It appears the missing messages caused the filter reset versus an algorithm problem. Also, the data-link hardware was designed for line-of-sight operations only. For these reasons, the final summary table includes lines to show what the floating filter, minimum indicator and MMAE errors would be without the filter reset.

The LAMBDA function returned a total of 29 unique ambiguity sets during the data run. The correct set was the top candidate returned from the first LAMBDA pair. The minimum indicator correctly chose the true ambiguity set at 4 minutes. It did not change except after the floating filter reset. The true ambiguity set was rediscovered at 15.16 minutes, after two other incorrect ambiguity sets were chosen.

After the filter reset, the true ambiguity set was returned from LAMBDA as the number one set, 2732 of 2876 epochs (95.0 percent), with a maximum ratio of 33.1.

Given the filter reset, the errors in the floating filter, minimum indicator and MMAE were larger than other data runs without the filter reset. If a filter reset occurred, the entire system was reset to include the floating filter, MMAE, and minimum indicator. Basically, the system was re-initialized. Figures A.47, A.48, and A.49 shows the floating filter performance.

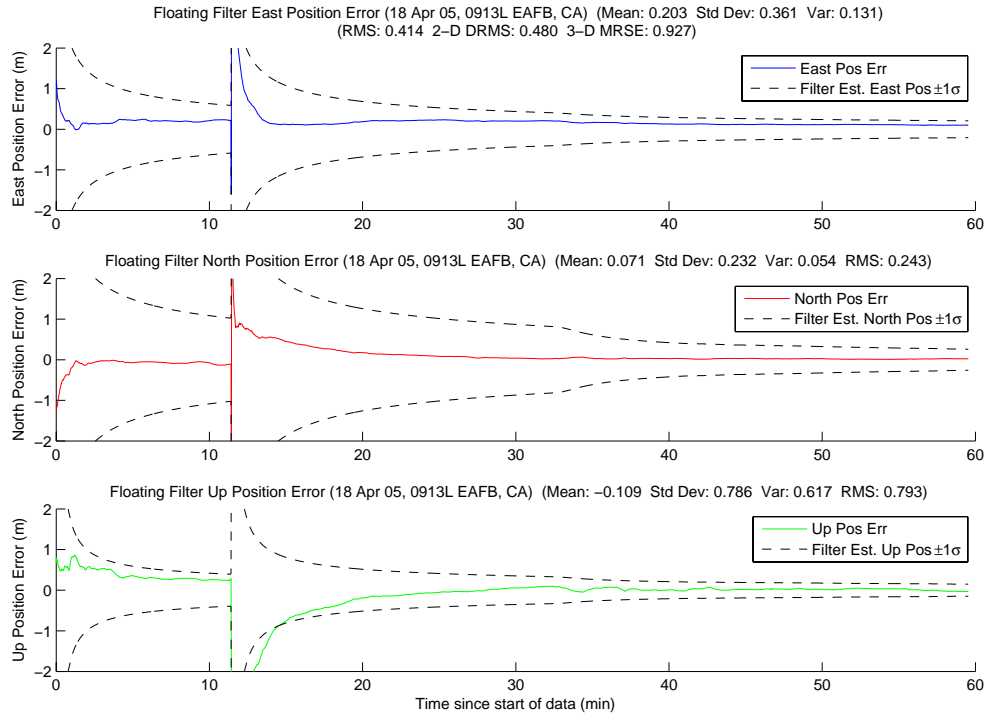


Figure A.47: Case F1.1: Floating Filter Position Errors

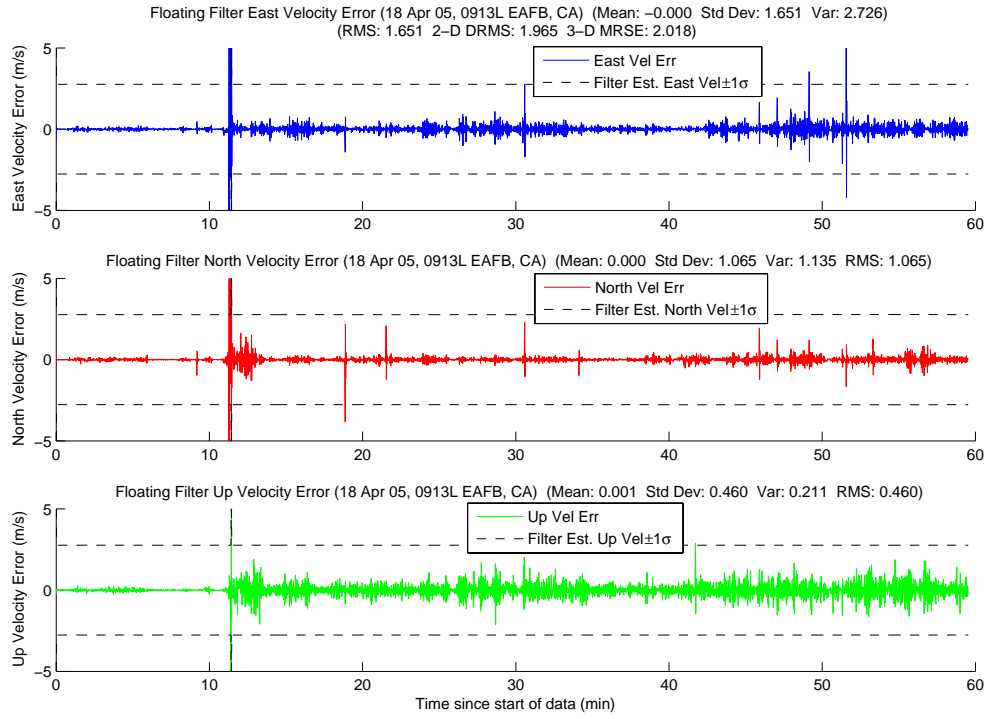


Figure A.48: Case F1.1: Floating Filter Velocity Errors

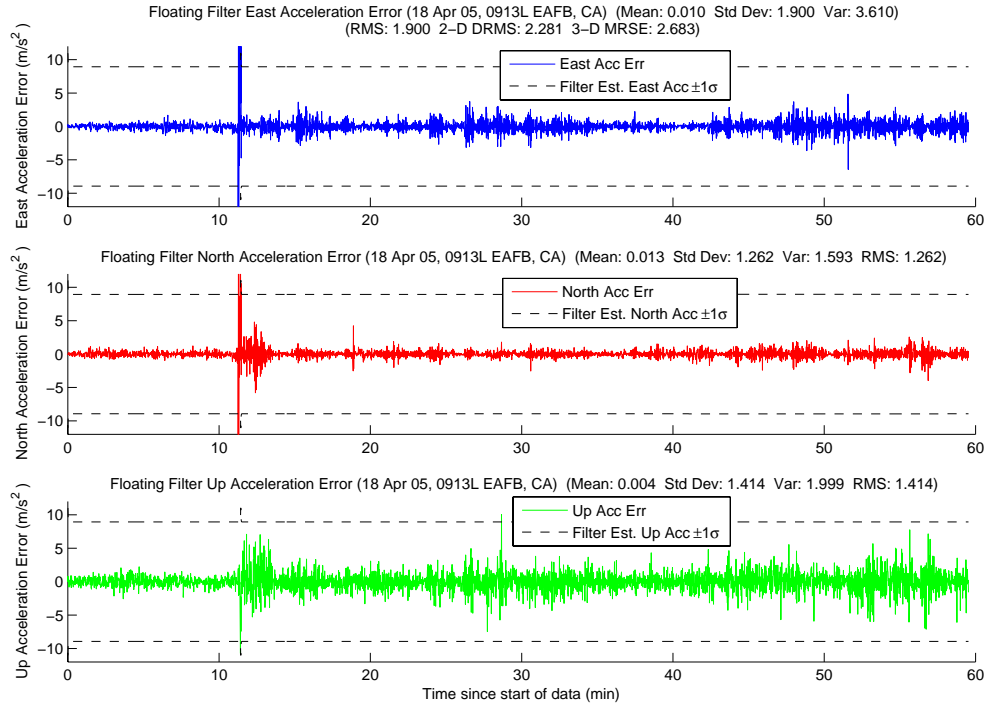


Figure A.49: Case F1.1: Floating Filter Acceleration Errors

Figure A.50 displays PRN 6 measurement information and the floating filter ambiguity estimate. The extremely large phase residuals at approximately 11 minutes coincide with the filter reset. The static initialization and other times when both aircraft were stopped can be seen in the phase residual plot at zero, seven and ten minutes. The rest of the satellite's plots follow in Figures A.51, A.52, A.53, A.54, A.55, A.56, and A.57.

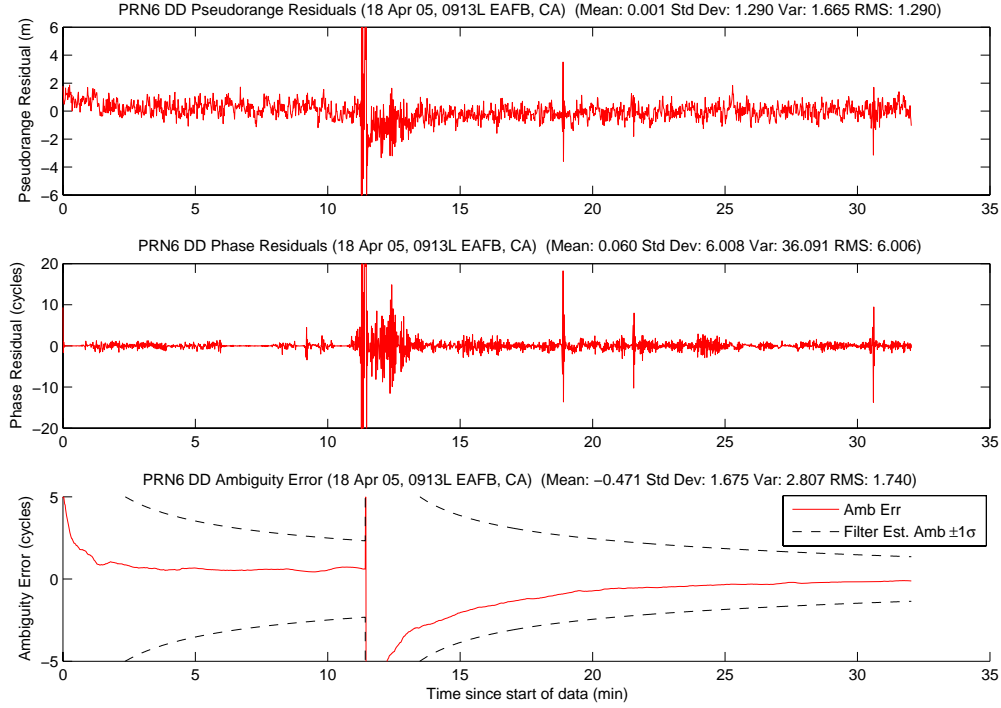


Figure A.50: Case F1.1: Satellite 6 Measurements

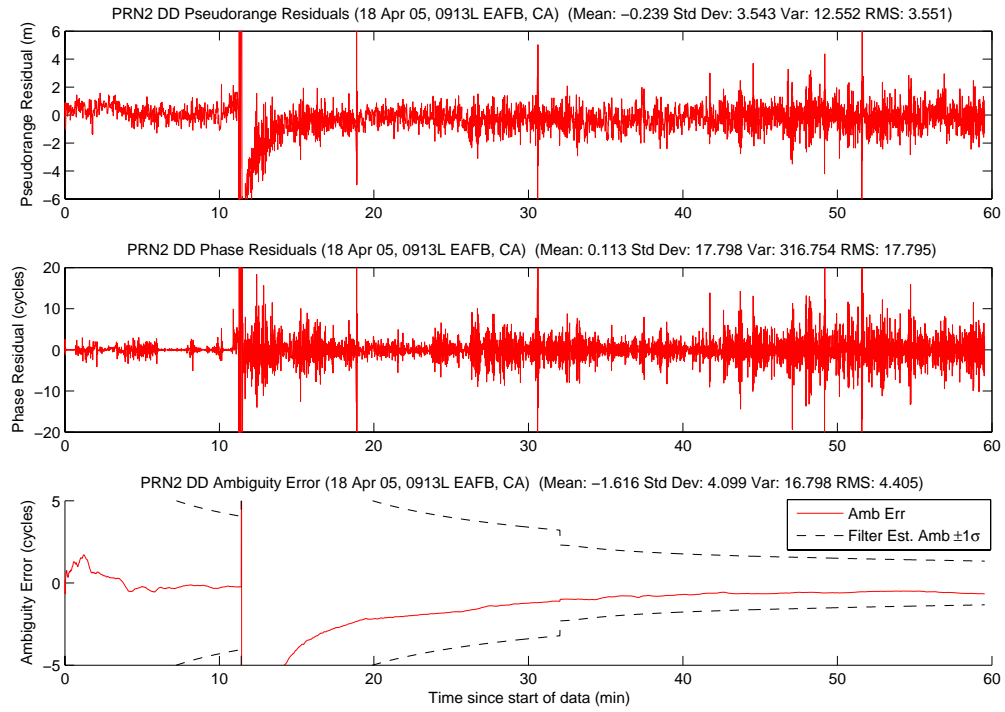


Figure A.51: Case F1.1: Satellite 2 Measurements

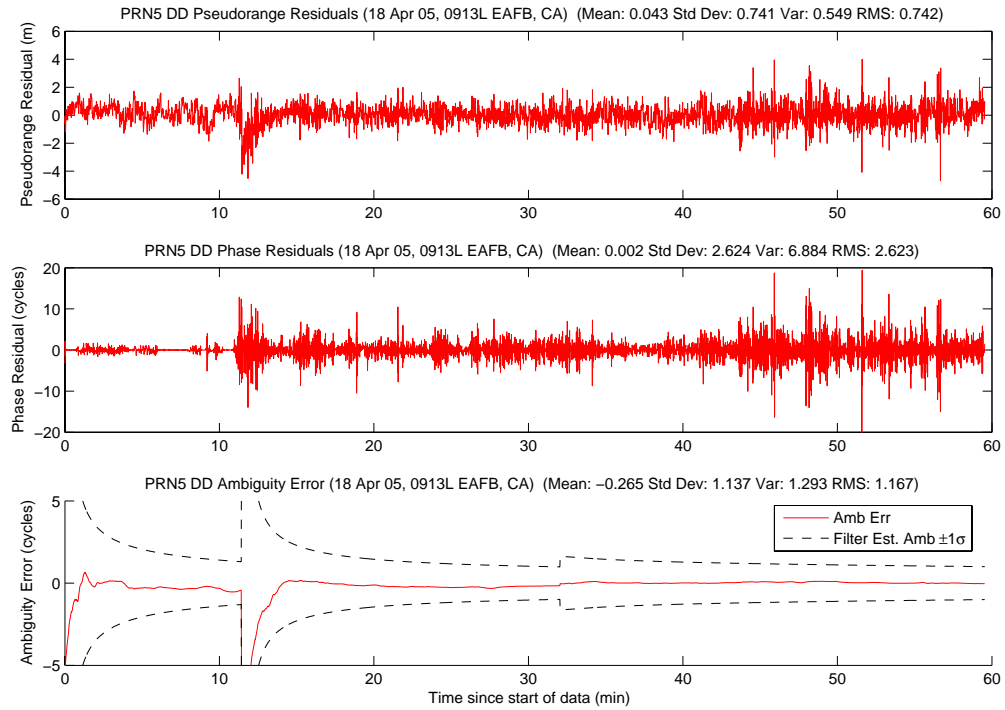


Figure A.52: Case F1.1: Satellite 5 Measurements

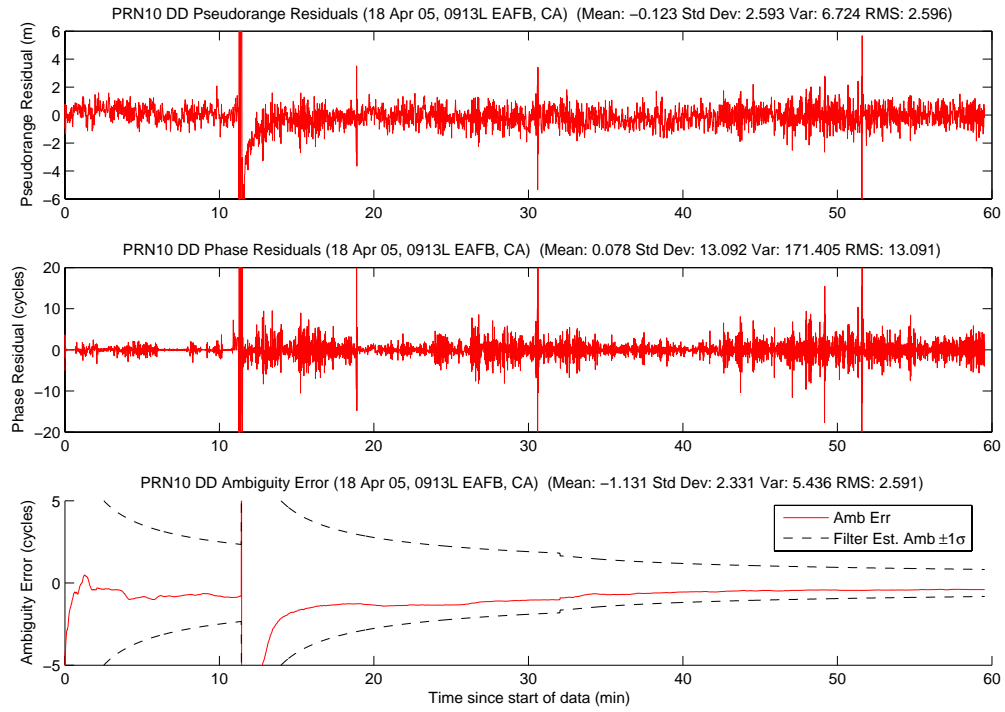


Figure A.53: Case F1.1: Satellite 10 Measurements

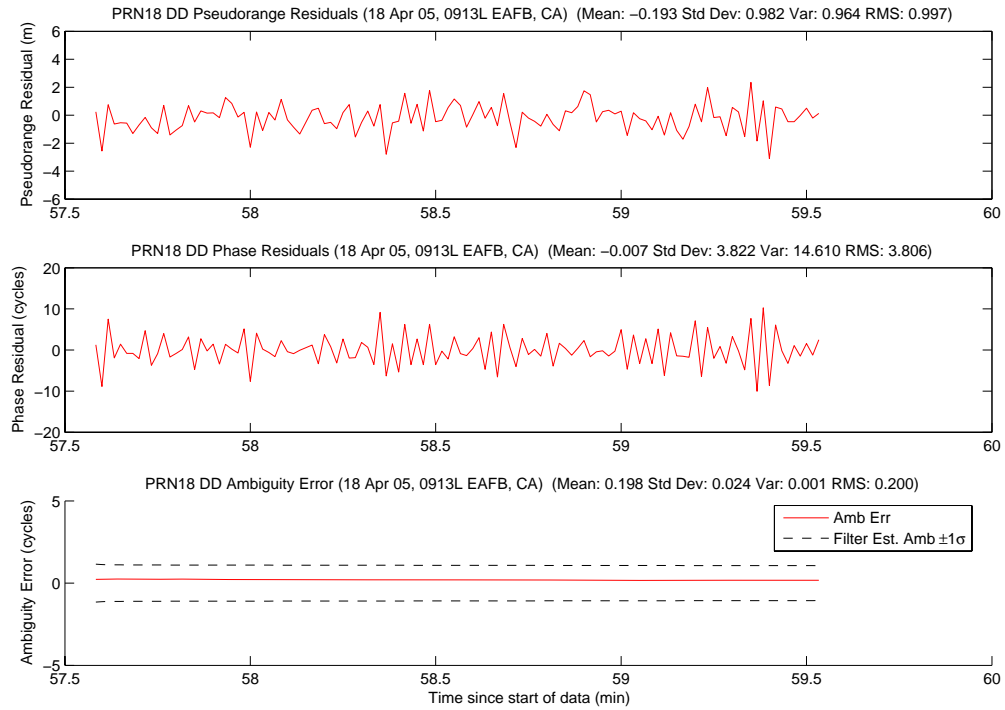


Figure A.54: Case F1.1: Satellite 18 Measurements

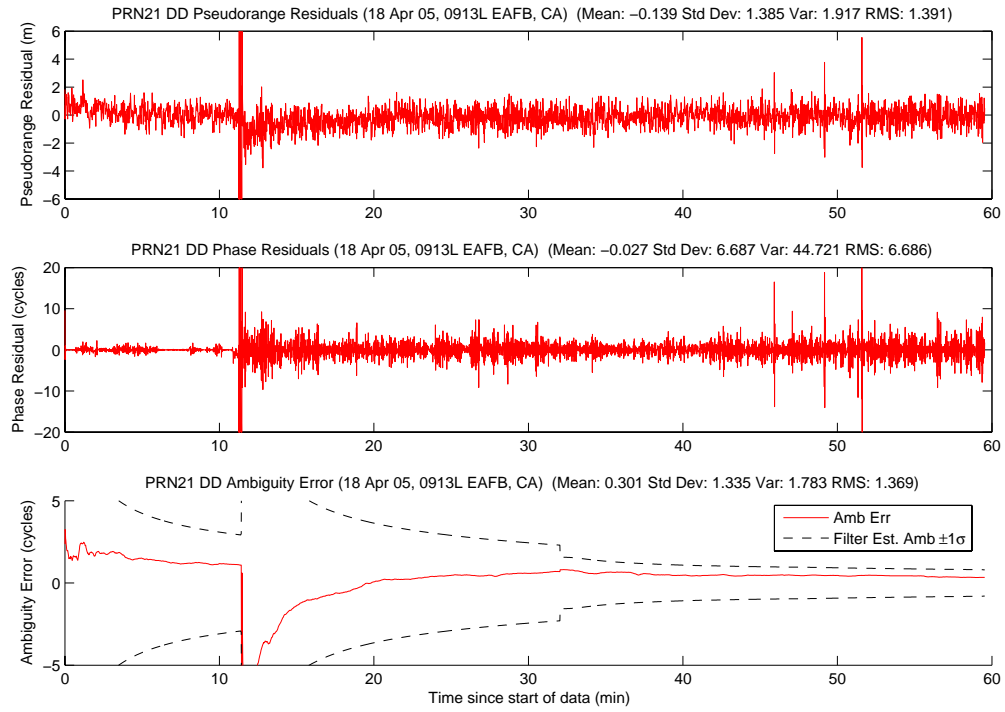


Figure A.55: Case F1.1: Satellite 21 Measurements

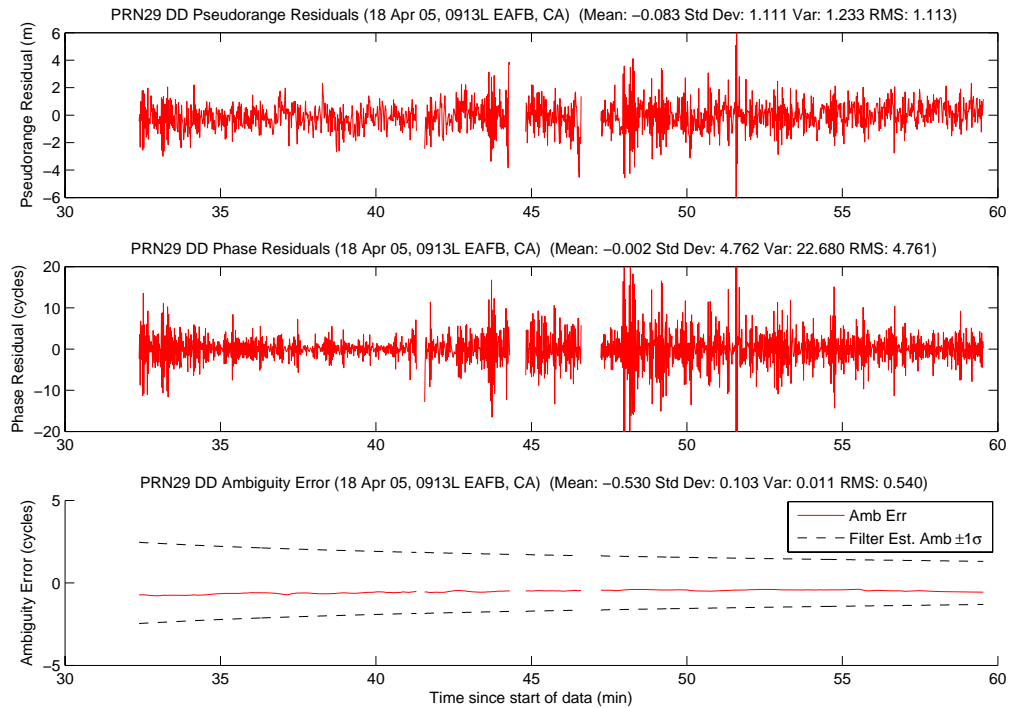


Figure A.56: Case F1.1: Satellite 29 Measurements

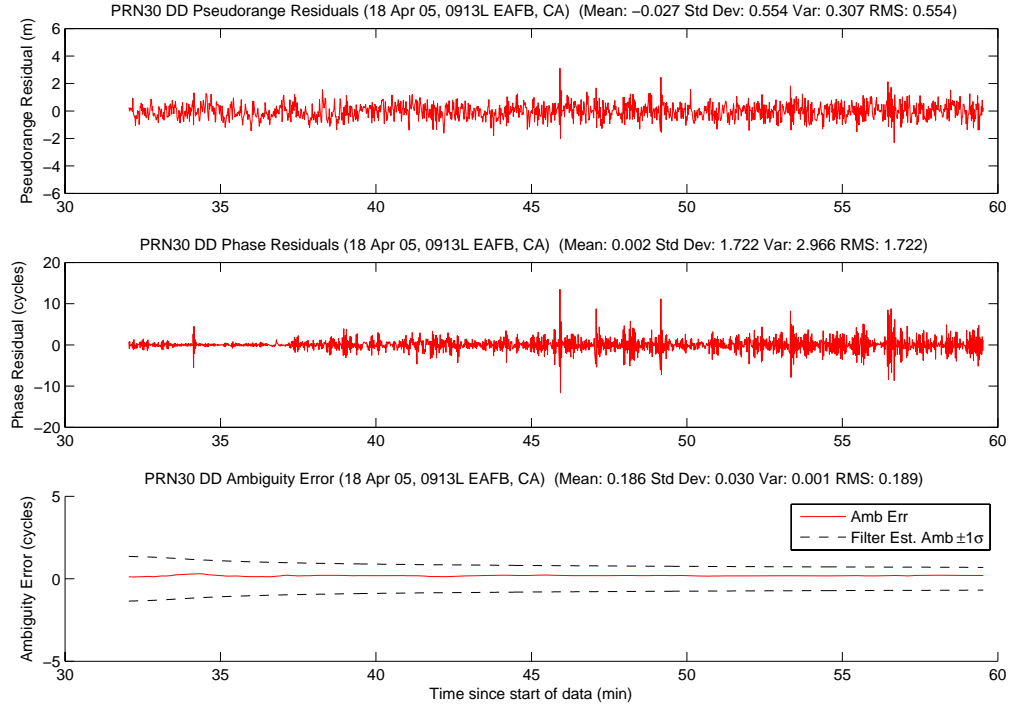


Figure A.57: Case F1.1: Satellite 30 Measurements

Figure A.58 shows the number of visible satellites. Most satellite dropouts were due to aircraft maneuvering. The C-12C had high mounted horizontal stabilizers and the GPS antenna was mounted on the center of the tail. When the aircraft banked, the horizontal stabilizer could mask certain satellites depending on the satellite geometry and bank angle. Since the elevation cutoff angle for satellites was 10 degrees, bank angles as small as 10 degrees could potentially mask satellites. It was not uncommon for the aircraft to bank 30 degrees or more during a data run. Even though the filter initialized with six satellites, the aircraft were static on the ramp with engines running, and the floating filter was able to converge normally.

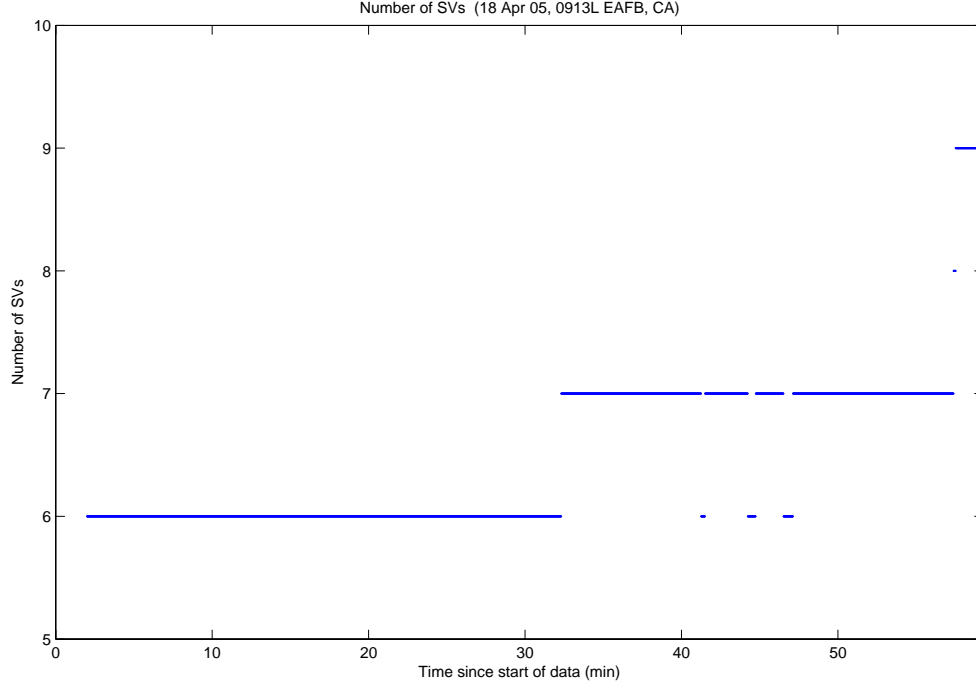


Figure A.58: Case F1.1: Visible Satellites

Once more, the residual plot, Figure A.59, is expanded to help identify the true ambiguity set's residuals. The filter reset at 11 minutes is clearly visible with a “wall” of erroneous ambiguity set's residuals. The probability plot, Figure A.60, has one point of interest. The filter reset occurred at 11 minutes when all the probabilities simultaneously drop to zero. Interestingly, the original EF[1] and what eventually became EF[15] have matching ambiguity sets. As the LAMBDA algorithm repopulated the ambiguity bank, an incorrect set initially absorbed 50 percent of the probability. Once the true set was returned from the LAMBDA function, it absorbed all of the probability by 26 minutes.

The minimum indicator discovered the true set, the second time, at 15.8 minutes (after selecting two wrong sets). The precise time the true ambiguity set was returned from the LAMBDA algorithm was not known. Assuming the minimum indicator immediately selected the true set after the initial 2 minute SOS residuals acquisition delay, the true ambiguity set was returned from LAMBDA at 13.8 minutes (approx-

imately 3 minutes after the reset). After the minimum indicator reselected the true set, it did not change during the remainder of the data run.

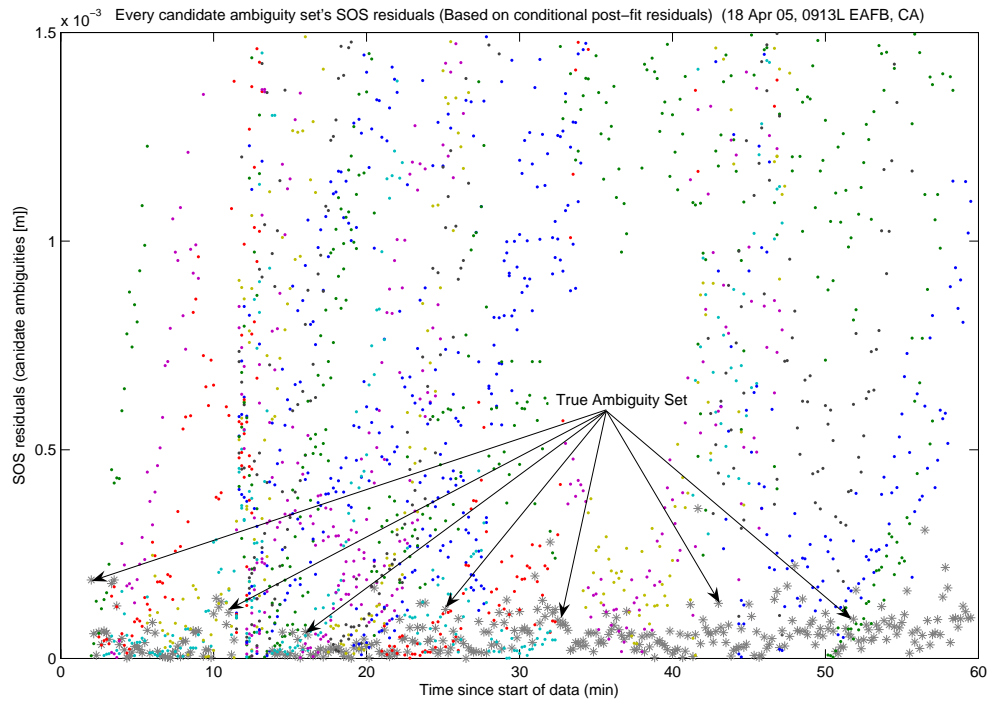


Figure A.59: Case F1.1: SOS Residuals

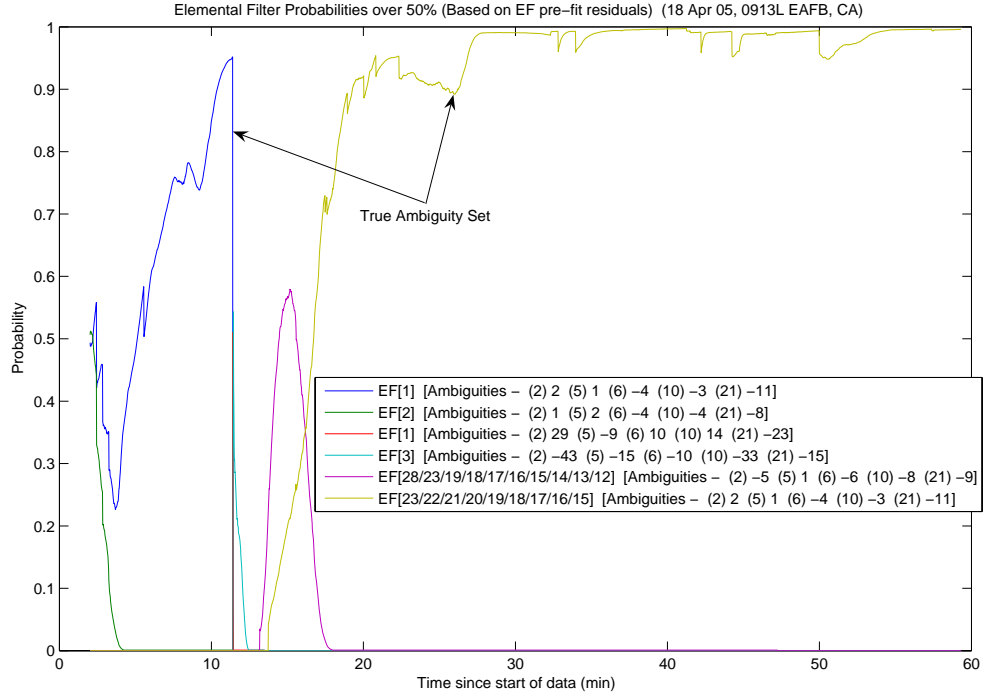


Figure A.60: Case F1.1: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.61 is EF[1], Figure A.62 is EF[2], and Figure A.63 is EF[23].

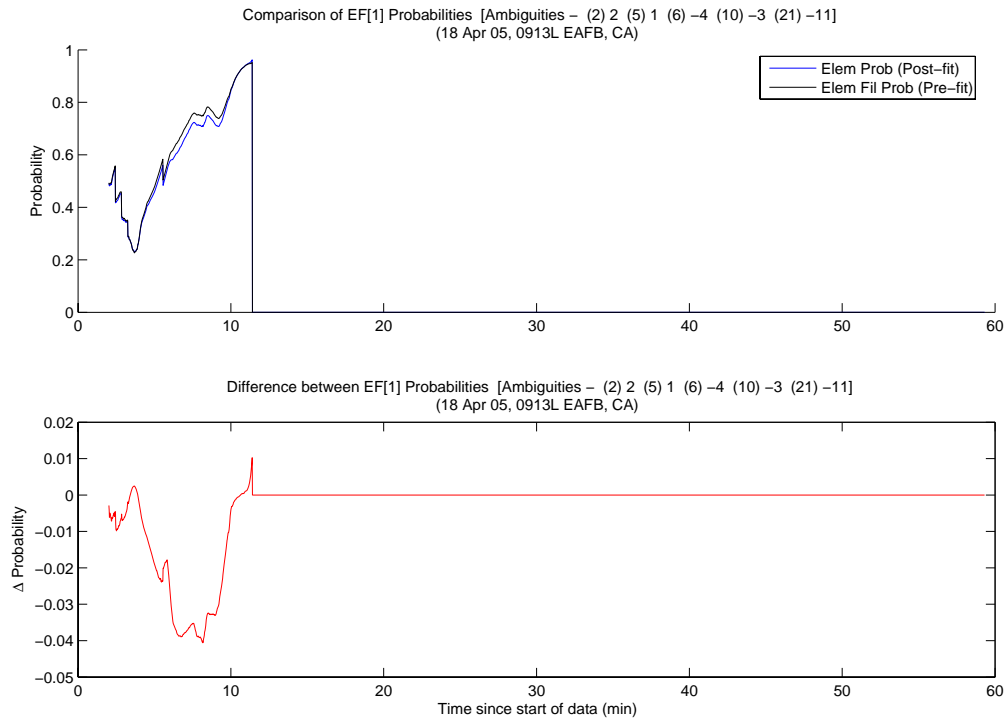


Figure A.61: Case F1.1: EF[1] Probability Comparison

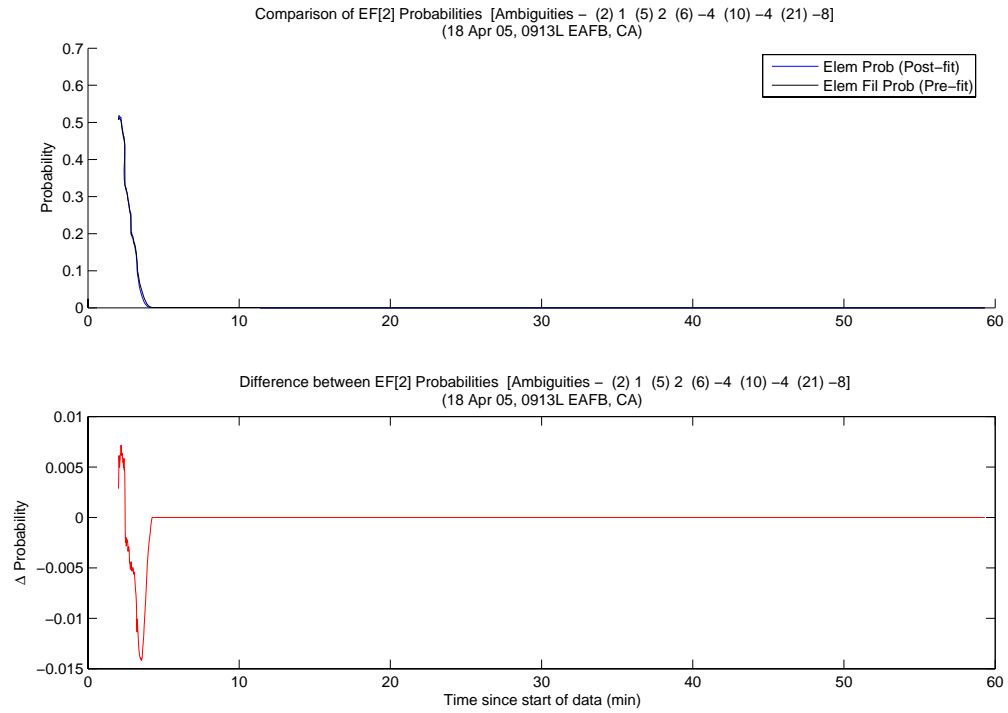


Figure A.62: Case F1.1: EF[2] Probability Comparison

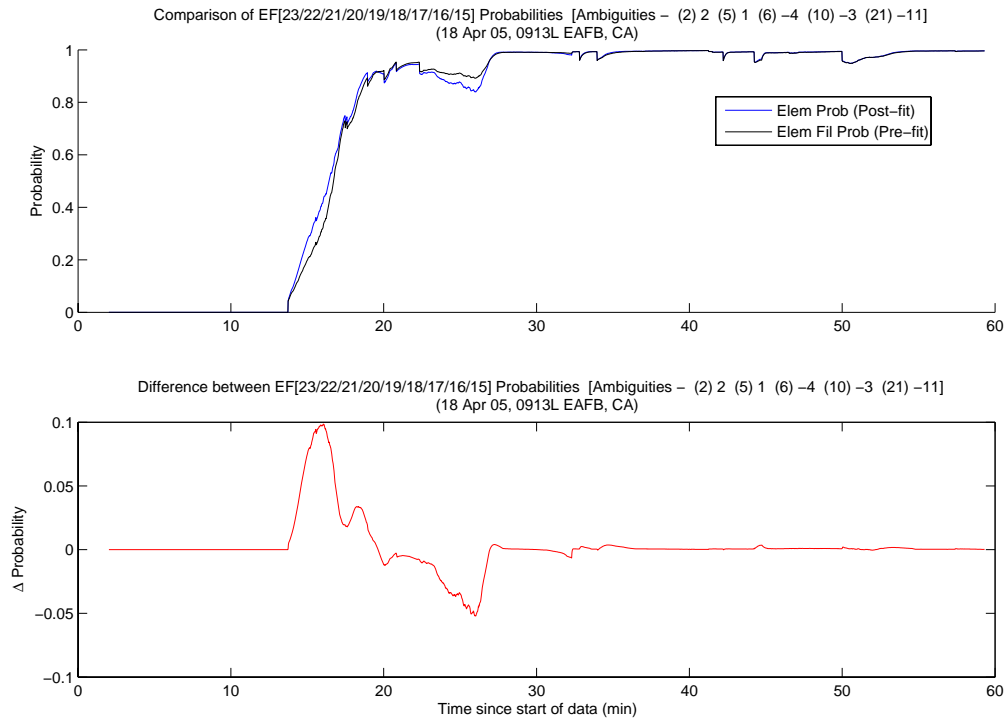


Figure A.63: Case F1.1: EF[23] Probability Comparison

The MMAE position error is shown in Figure A.64. The majority of the error was from the filter reset at 11 minutes. Smaller divergences in the MMAE position error due to probability fluctuations are not visible.

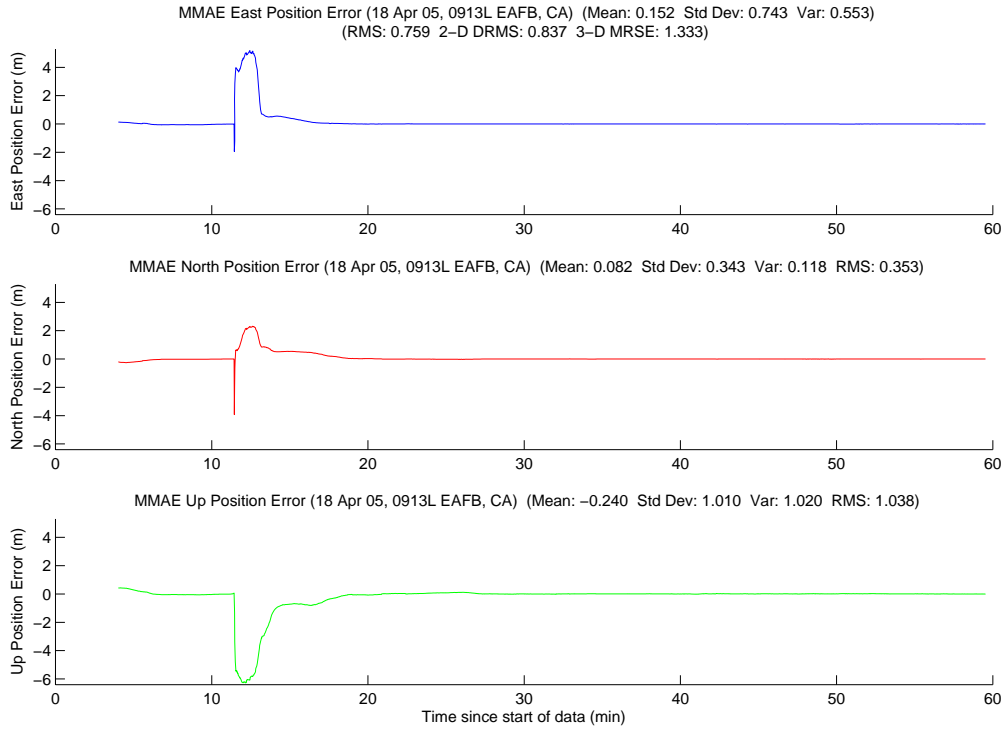


Figure A.64: Case F1.1: MMAE Position Error

Finally, Figure A.65 displays the position error plot for the minimum indicator variable, and Table A.5 gives a summary of each method's position error. Table A.5 also includes the post processing error numbers since the filter did not reset while in post processing. Just like the MMAE position errors, the minimum indicator error also suffered from the filter reset. The error immediately dropped at 15.8 minutes when the true ambiguity set was reacquired. There were no minimum indicator values from the floating filter reset at 11 minutes, until the first incorrect ambiguity set was selected at approximately 13 minutes. Therefore, there is a gap in the minimum indicator during those times.

Unfortunately, the minimum indicator error was not within the desired centimeter-level accuracy for MRSE with the floating filter reset. Of course, without the filter reset the minimum indicator remained well within the tolerance and was similar in magnitude to the dynamic ground test results. Most importantly, the algorithm and hardware performed extremely well during the flight where the system was designed

to run (i.e., close formation) . In every flight test run, when the aircraft were in close formation, the algorithm was within the centimeter-level requirement.

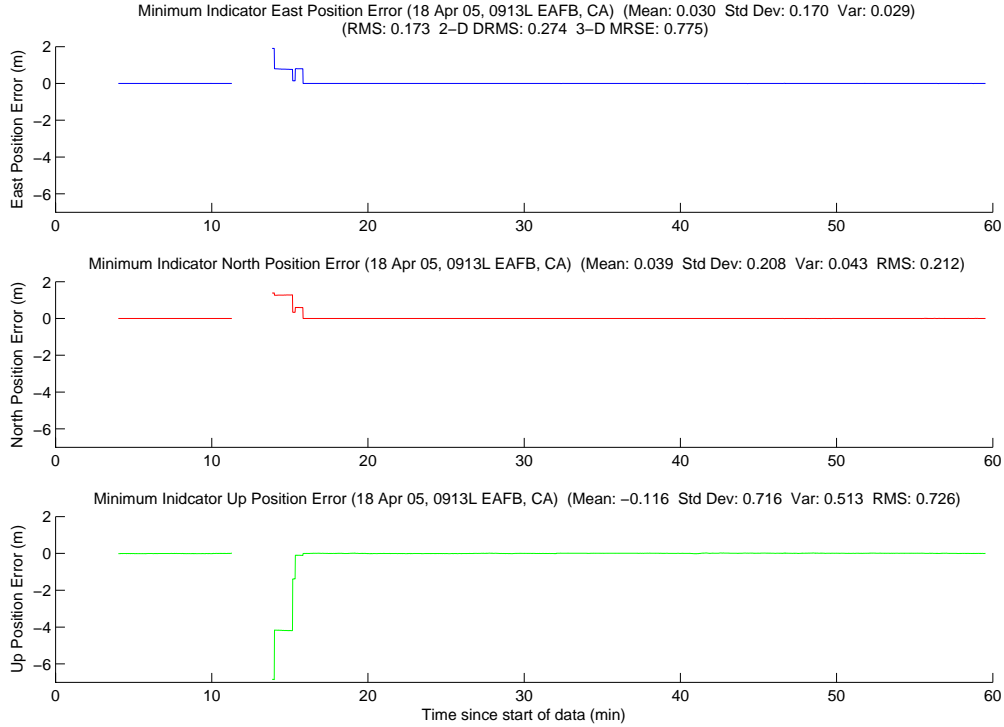


Figure A.65: Case F1.1: Minimum Indicator Position Error

Table A.5: Case F1.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.203	0.361	0.071	0.232	-0.109	0.786	0.480	0.927
MMAE	0.152	0.743	0.082	0.343	-0.240	1.010	0.837	1.333
Min. Ind.	0.030	0.170	0.039	0.208	-0.116	0.716	0.274	0.775
Float Filter*	0.165	0.060	-0.046	0.087	0.169	0.144	0.201	0.299
MMAE*	0.001	0.024	-0.010	0.037	0.020	0.065	0.045	0.082
Min. Ind.*	-0.000	0.001	-0.000	0.001	0.002	0.008	0.002	0.009

(Note: * are the values without the floating filter reset at 11 minutes.)

A.2.2 Case F1.2: Flight Test 1, Second Data Set. Case F1.2 starts while both aircraft are airborne approximately eight minutes after case F1.1 end. Once more, the data set is one hour long. During the test, the aircraft ranged from 34.4 to

2483.6 meters apart. Figure A.66 shows the minimum indicator North, East, Down relative positions.

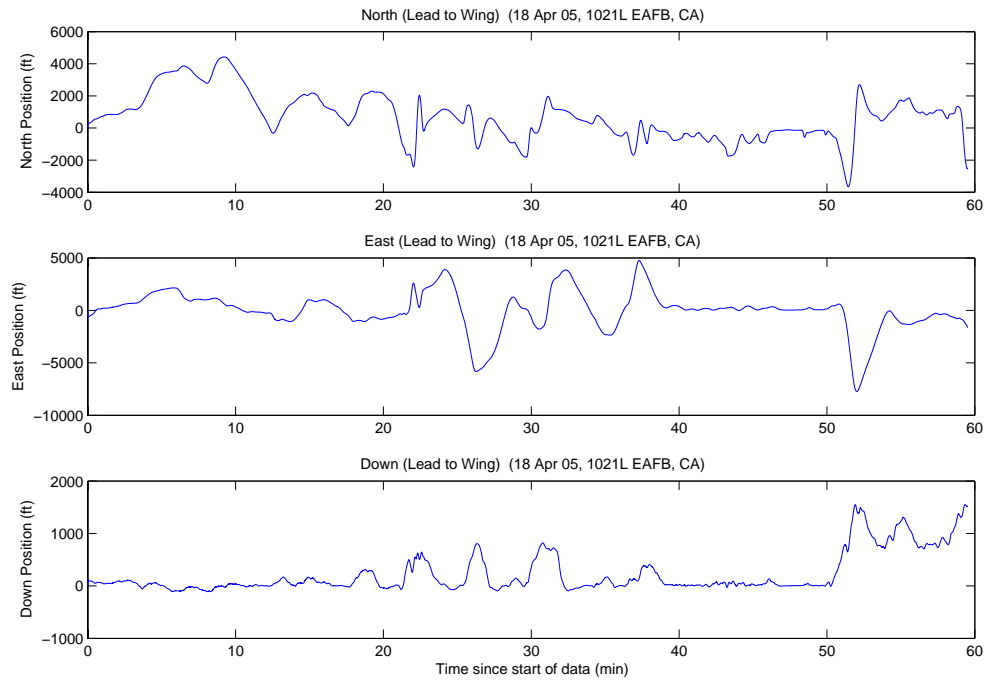


Figure A.66: Case F1.2: North, East, Down Relative Positions

There is one point worth mentioning before the rest of data is displayed. Some very aggressive maneuvering by both aircraft at 52 minutes into the data run caused multiple cycle slips. Both aircraft were banking over 45 degrees in either direction and pulling approximately 2.5 gs. Due to the cycle slips, the minimum indicator was wrong during the last eight minutes of the data run. There was not enough time for LAMBDA to return the true set before the end of the data run. More discussion about the true ambiguity set follows with the SOS residuals plot.

The LAMBDA function returned a total of 41 unique ambiguity sets during the data run. The correct set was the top candidate returned from the first LAMBDA pair. The minimum indicator correctly chose the true ambiguity set at 4 minutes. It did not change except once at 22.8 minutes when there were only five visible satellites. The minimum indicator incorrectly chose a set with PRN 10 ambiguity off by one cycle. The minimum indicator quickly changed back to the true set at 24.1 minutes. The true set was returned from LAMBDA as number one, 2966 of 3436 epochs (86.3 percent), with a maximum ratio of 53.4.

Figures A.67, A.68, and A.69 show the floating filter errors. The floating filter performance was generally very good, until approximately 52 minutes when the cycle slips occurred. The floating filter performed as expected with an average of nine visible satellites.

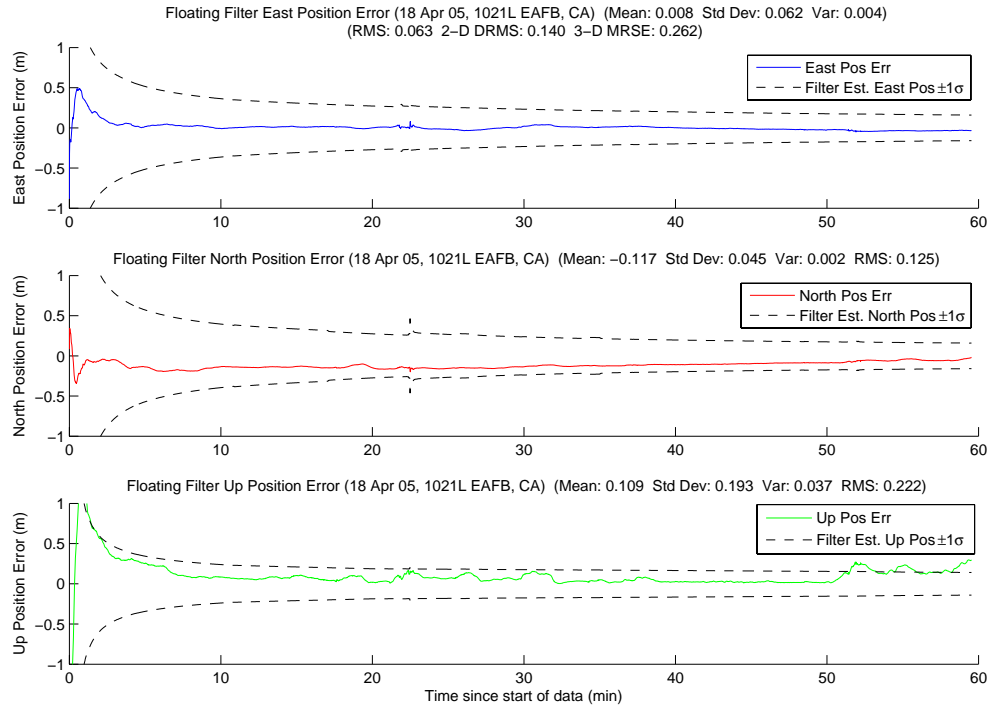


Figure A.67: Case F1.2: Floating Filter Position Errors

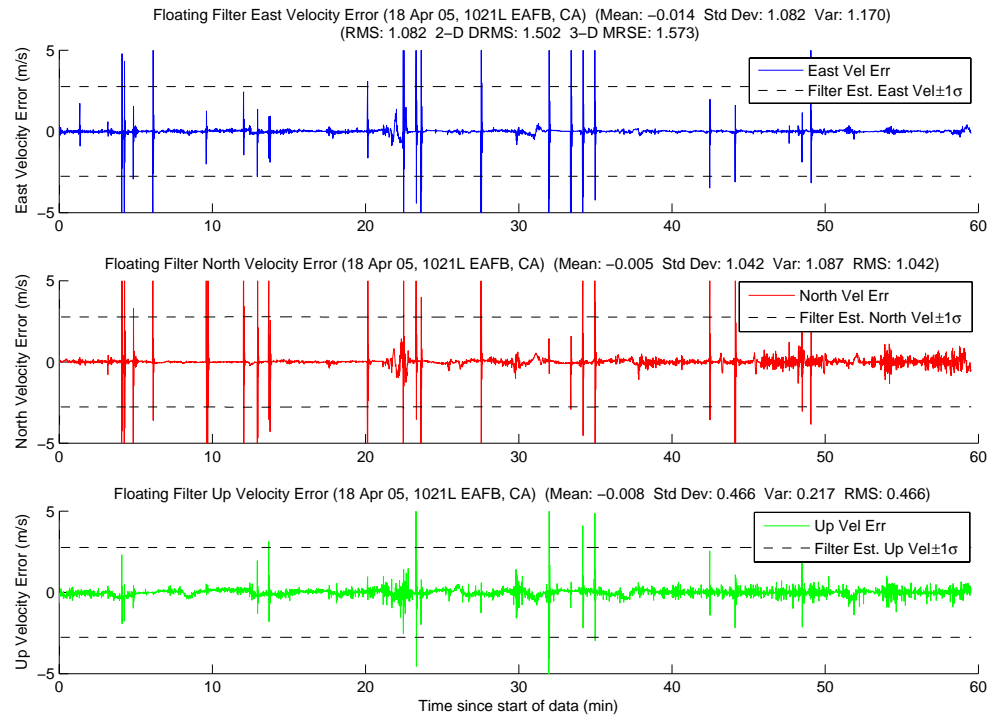


Figure A.68: Case F1.2: Floating Filter Velocity Errors

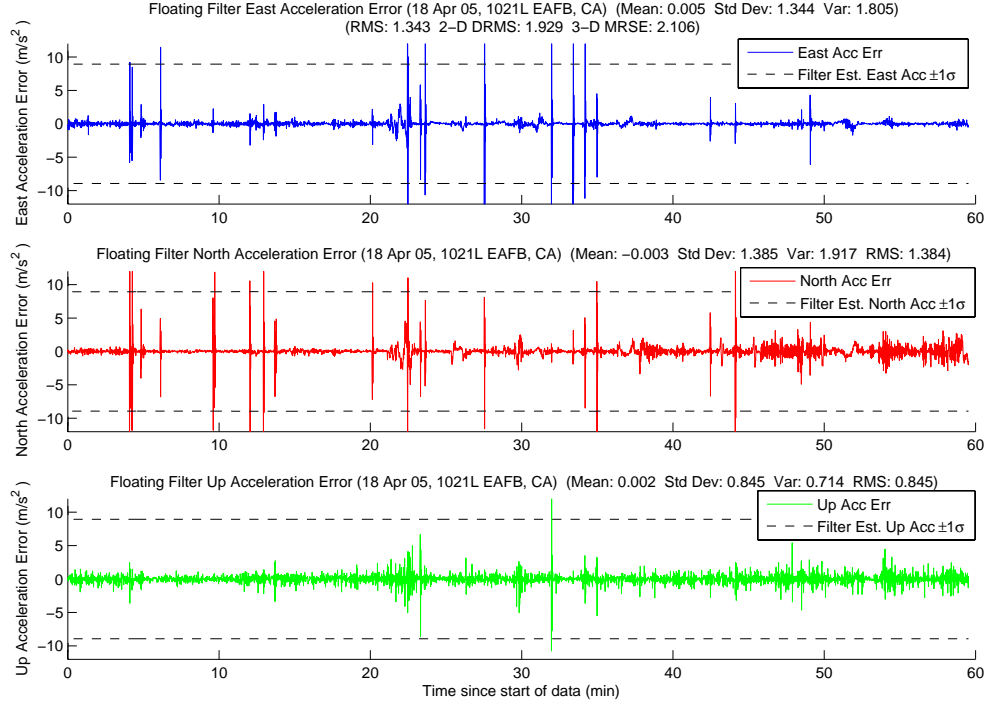


Figure A.69: Case F1.2: Floating Filter Acceleration Errors

Figure A.70 displays PRN 15 measurement information and the floating filter ambiguity estimate. PRN 15 was one of the satellites that experienced a cycle slip, as seen in the small bumps in the filter estimated 1σ values for the ambiguity error at 52 minutes. The rest of the satellite's plots follow in Figures A.71, A.72, A.73, A.74, A.75, A.76, A.77, A.78, A.79, and A.80.

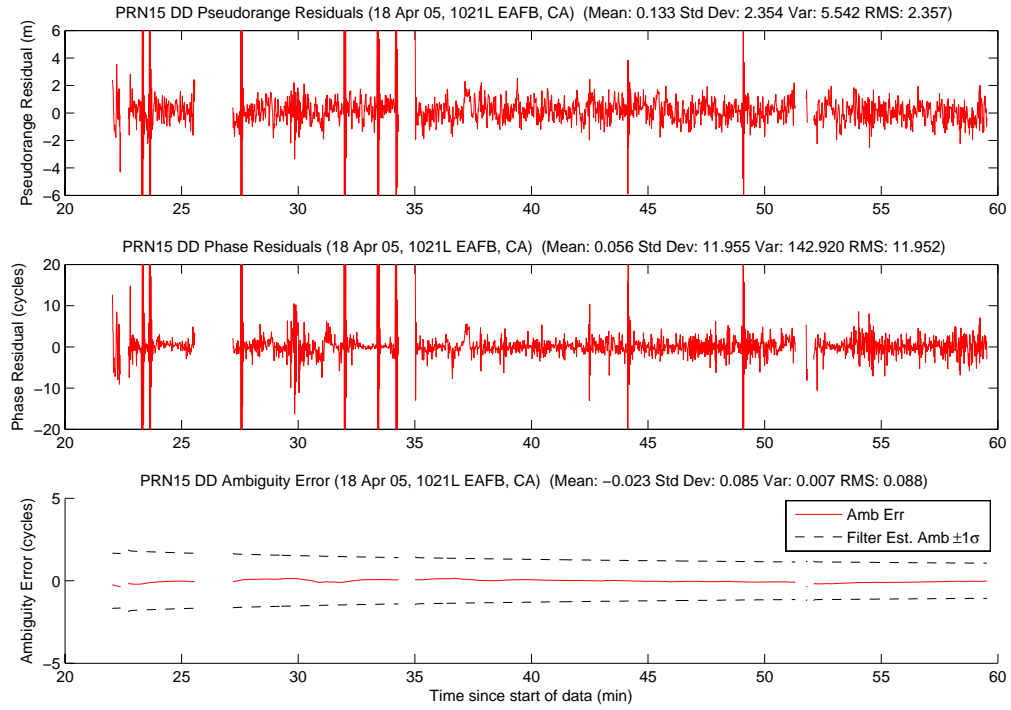


Figure A.70: Case F1.2: Satellite 15 Measurements

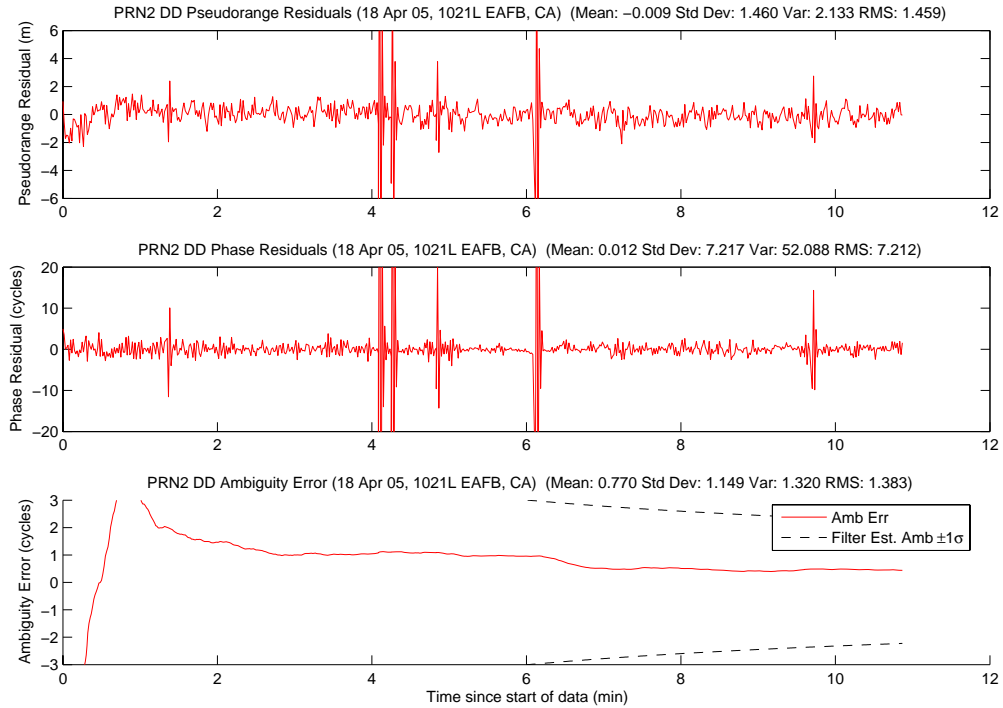


Figure A.71: Case F1.2: Satellite 2 Measurements

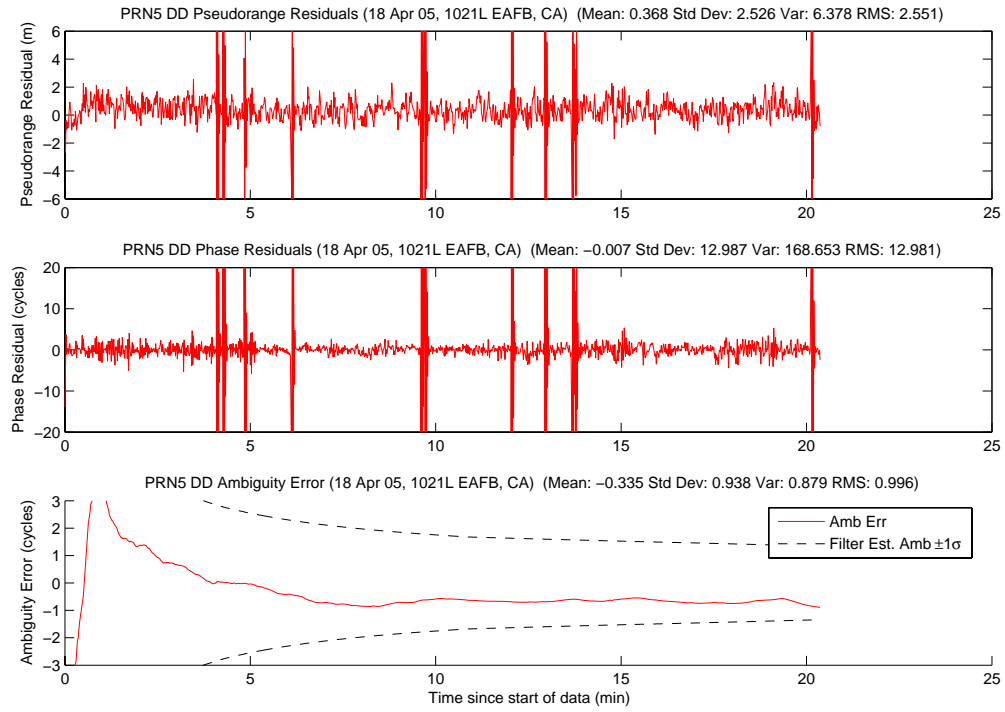


Figure A.72: Case F1.2: Satellite 5 Measurements

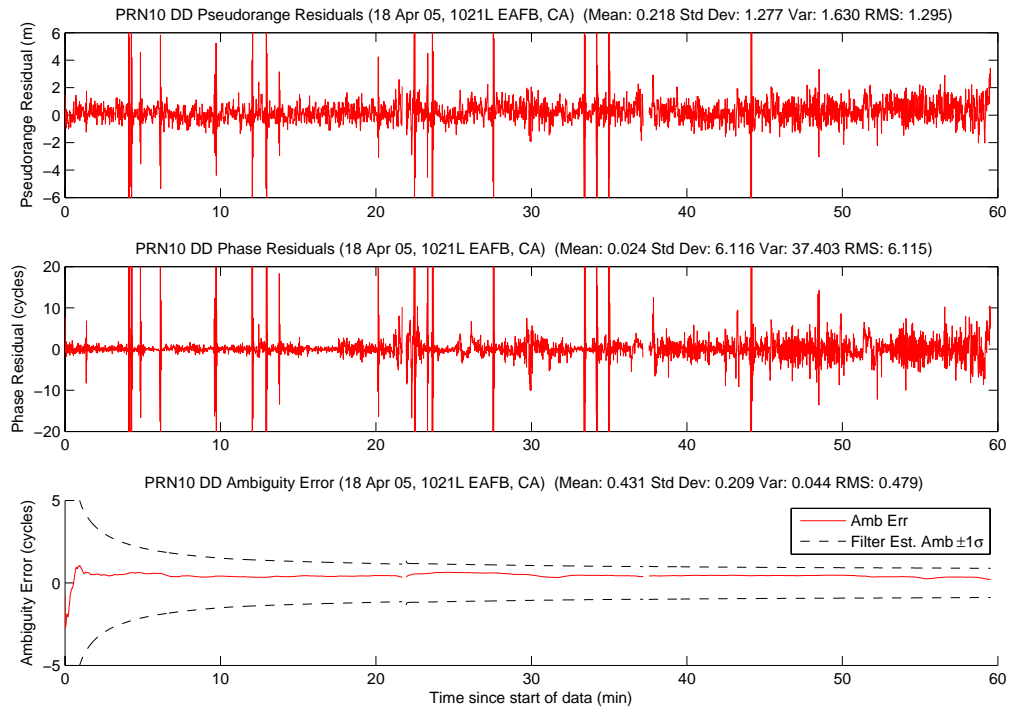


Figure A.73: Case F1.2: Satellite 10 Measurements

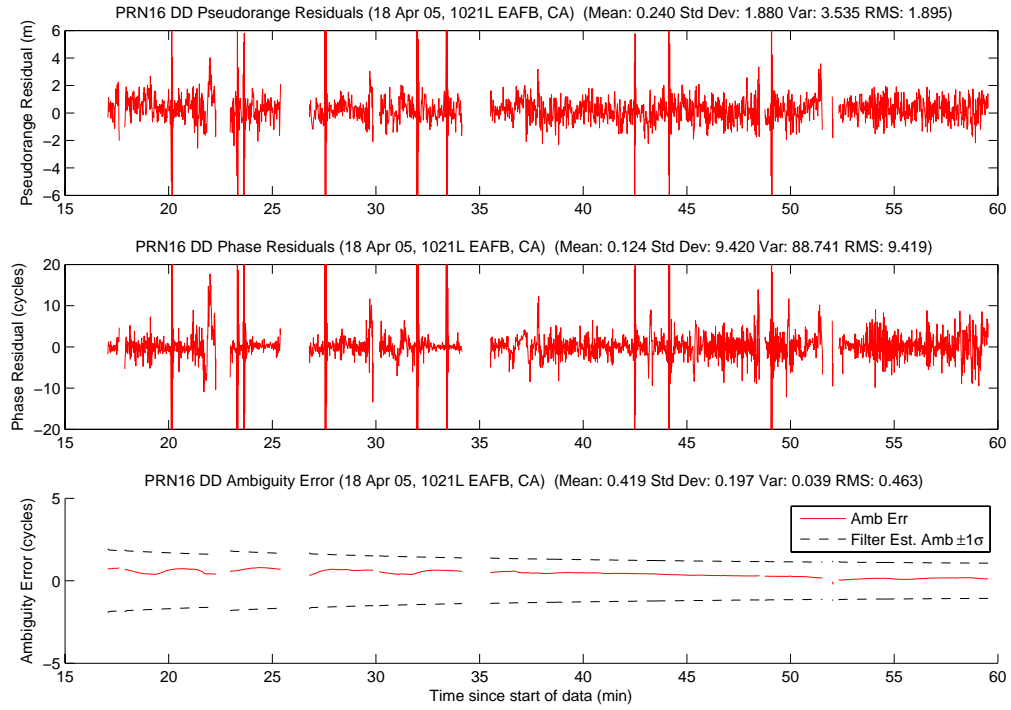


Figure A.74: Case F1.2: Satellite 16 Measurements

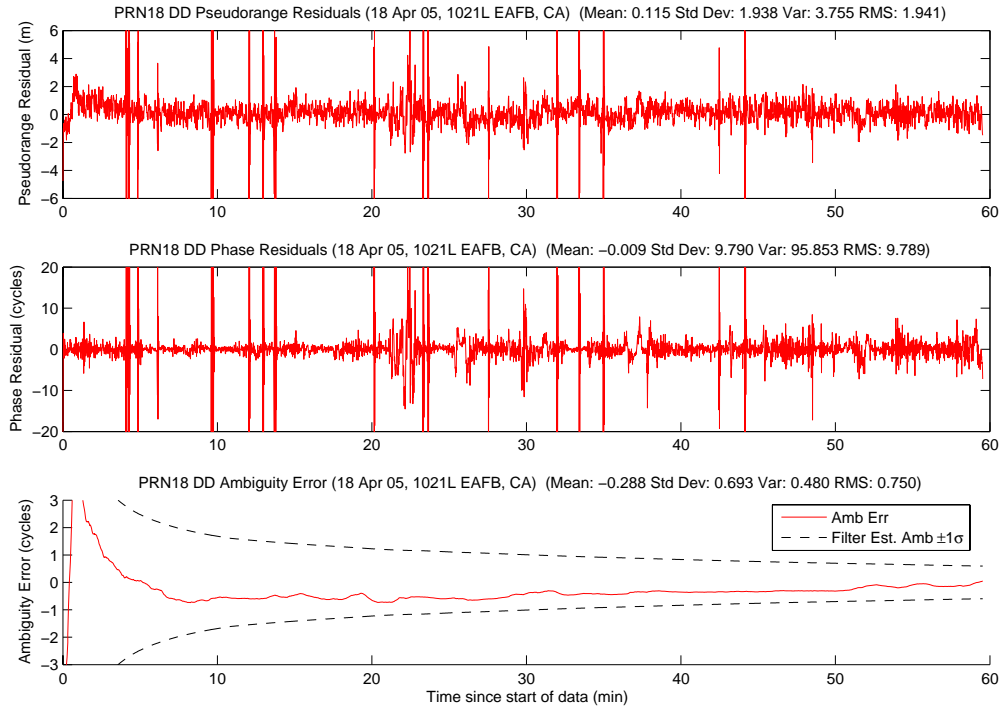


Figure A.75: Case F1.2: Satellite 18 Measurements

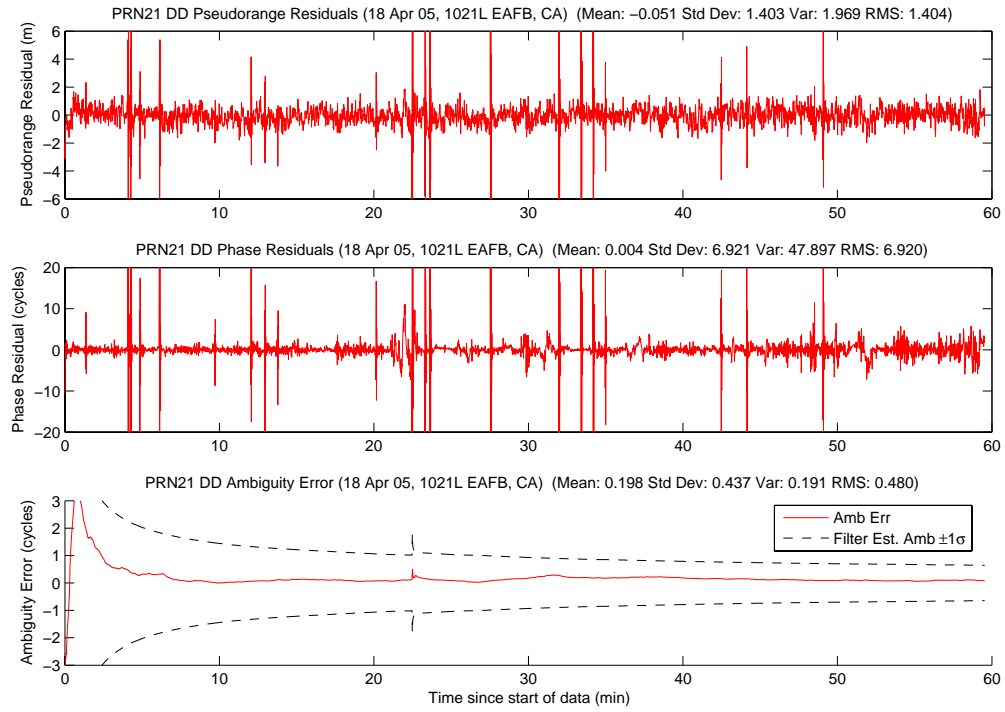


Figure A.76: Case F1.2: Satellite 21 Measurements

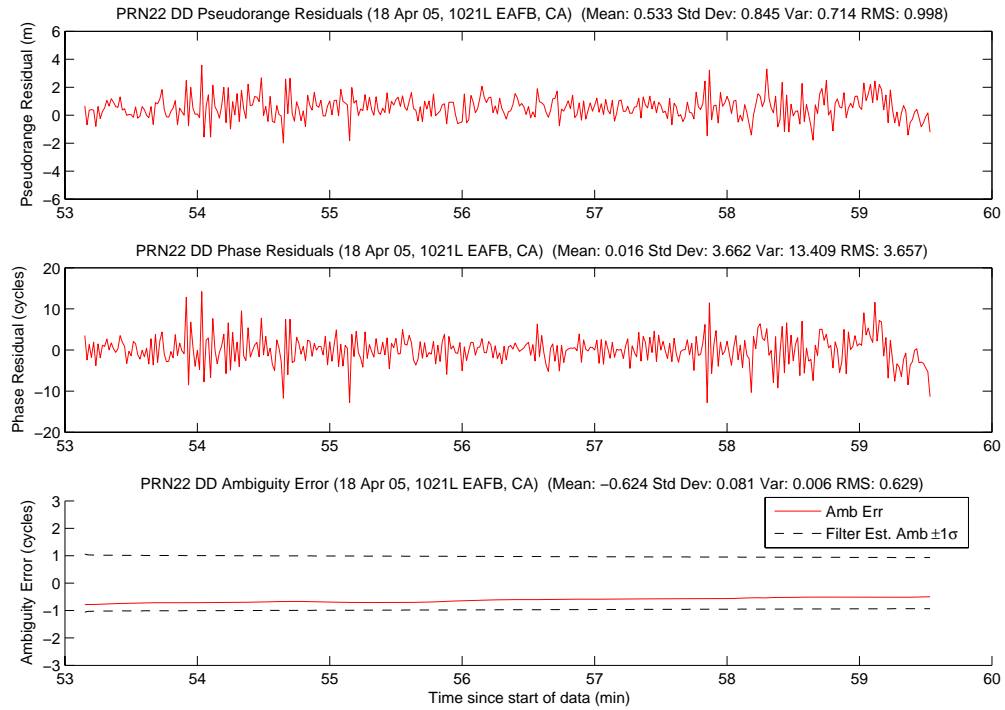


Figure A.77: Case F1.2: Satellite 22 Measurements

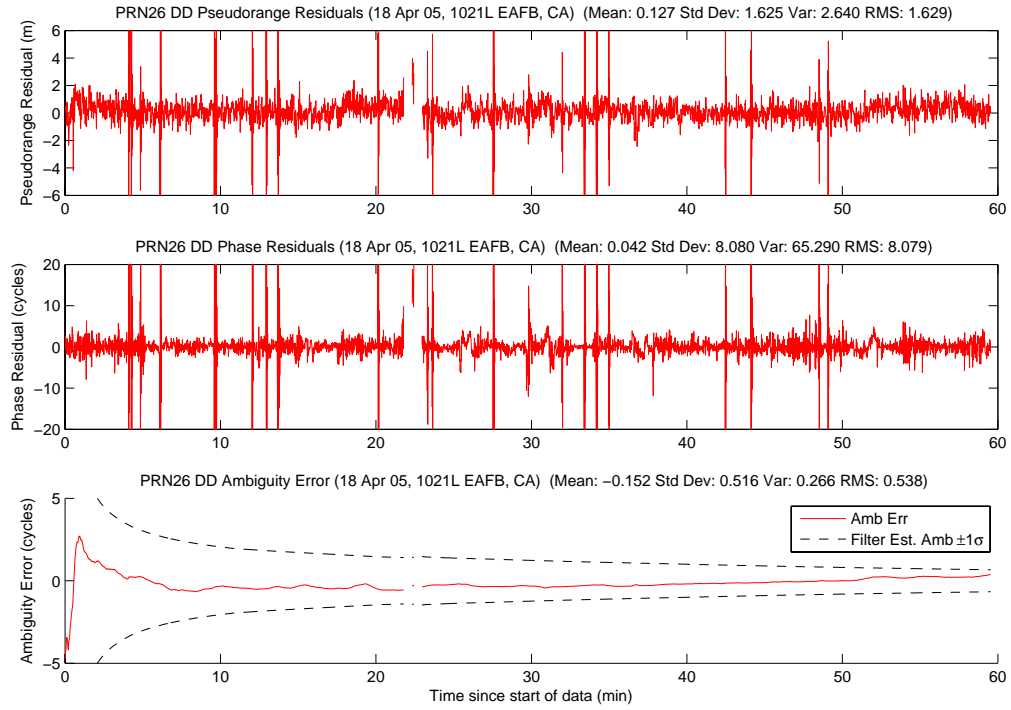


Figure A.78: Case F1.2: Satellite 26 Measurements

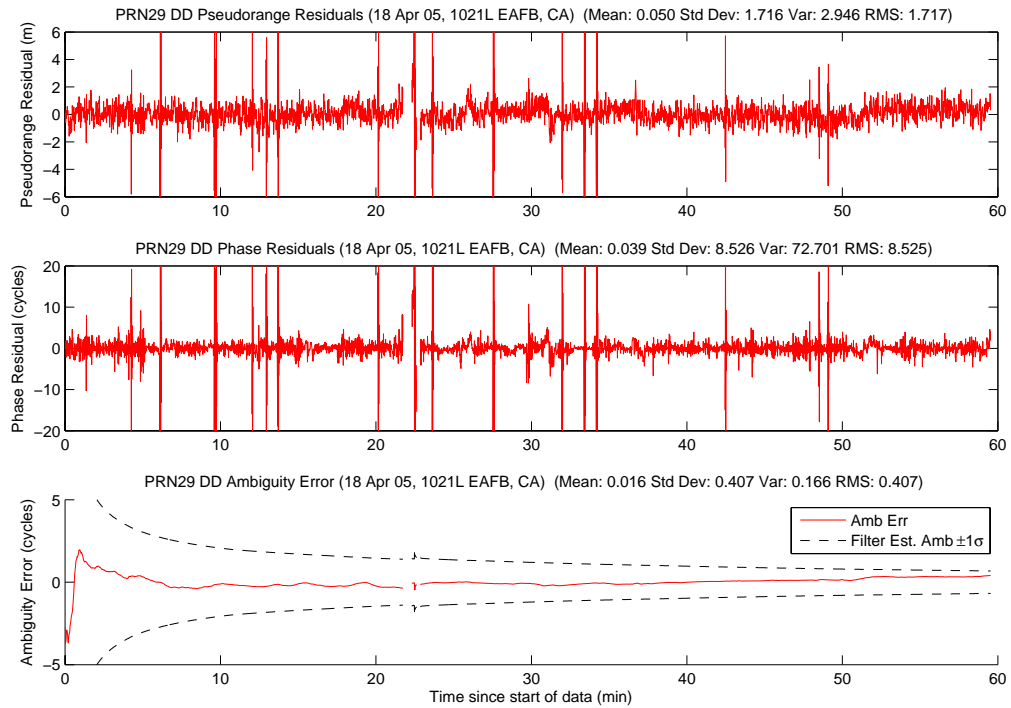


Figure A.79: Case F1.2: Satellite 29 Measurements

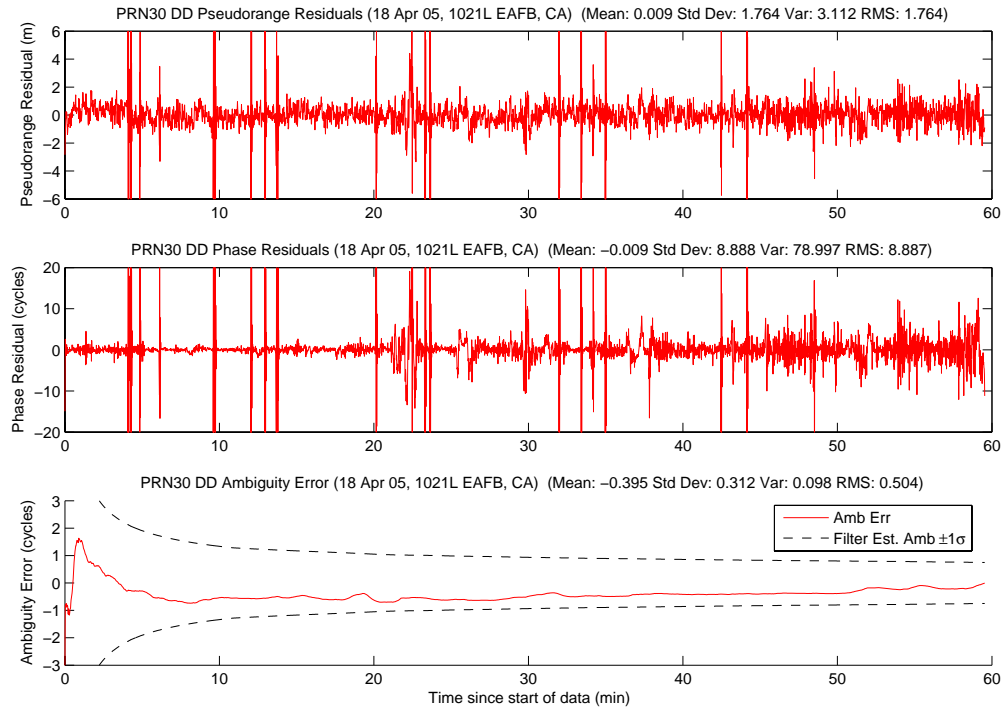


Figure A.80: Case F1.2: Satellite 30 Measurements

Figure A.81 shows the number of visible satellites. All of the satellite dropouts prior to 52 minutes were due to aircraft maneuvering. With the low elevation cutoff angle of ten degrees used during flight testing and the high mounted horizontal stabilizer, any aircraft bank angle over ten degrees had the potential to mask a satellite. Since the aircraft were operating at altitudes over 10,000 feet AGL, rising and setting satellites were visible at elevations below the horizon. A lower elevation cutoff angle could potentially be used for aircraft with lower mounted horizontal stabilizers, or for very benign flying conditions.

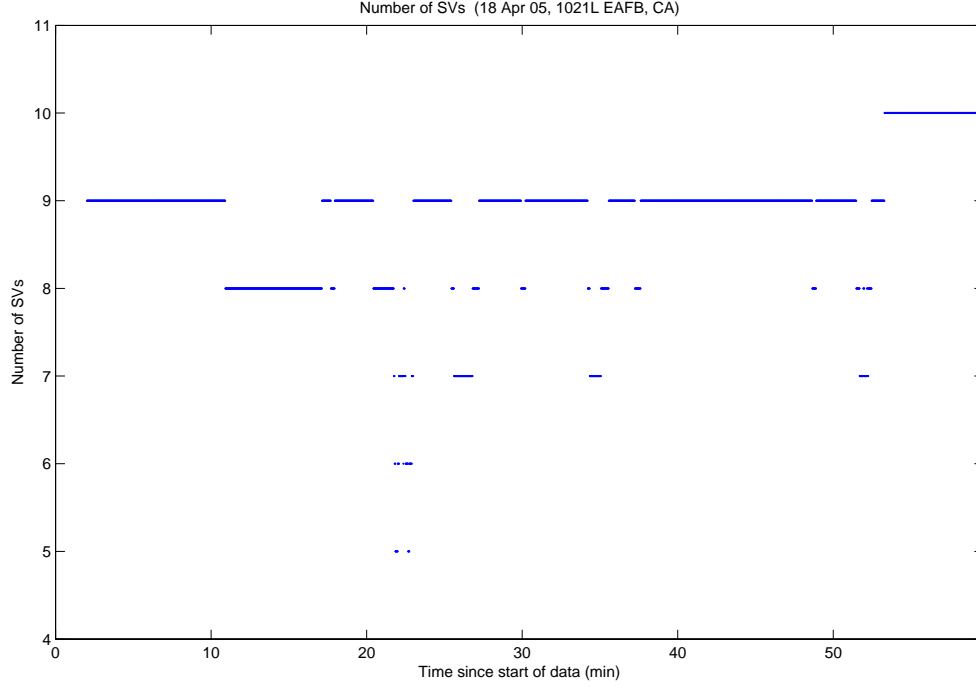


Figure A.81: Case F1.2: Visible Satellites

The SOS residuals plot, Figure A.82, shows some interesting results. The true ambiguity set's SOS residuals were easily detectable up until the cycle slip at 52 minutes. After the cycle slips, there was no set which appeared to “fill the gap” left by the true ambiguity set. This was the tale-tell sign that the true ambiguity set was not in the ambiguity bank. The probability plot, Figure A.83, reveals the same information. Notice the numerous ambiguity sets which absorbed over 99 percent of the conditional probability during the last eight minutes of the run. Since the true ambiguity set was not in the ambiguity bank, neither the minimum indicator nor the MMAE could correctly identify the true set.

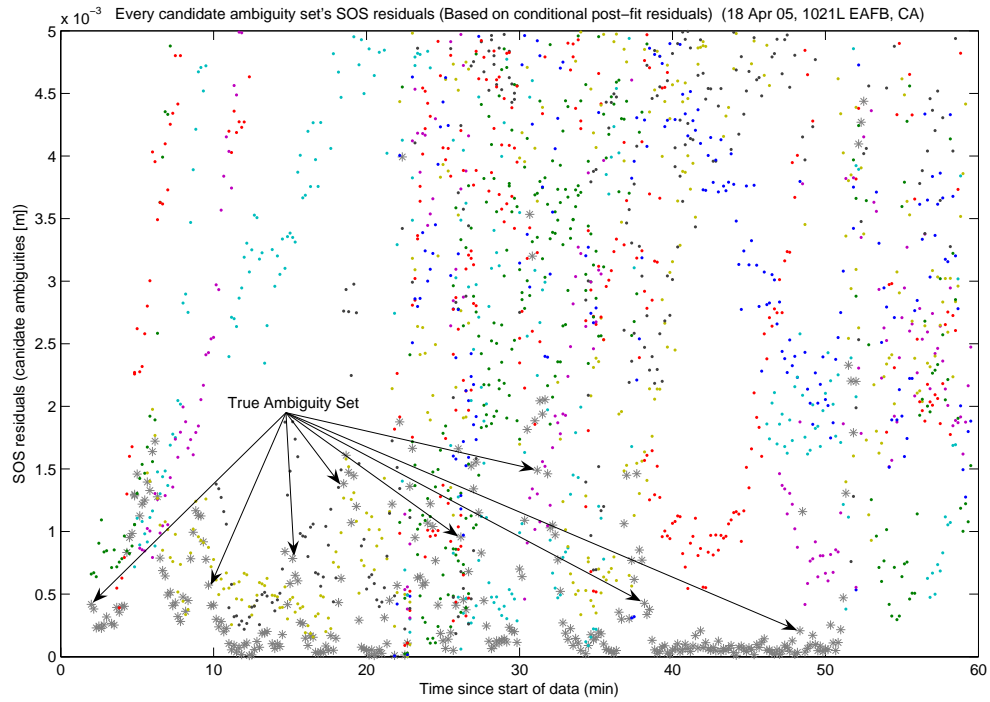


Figure A.82: Case F1.2: SOS Residuals

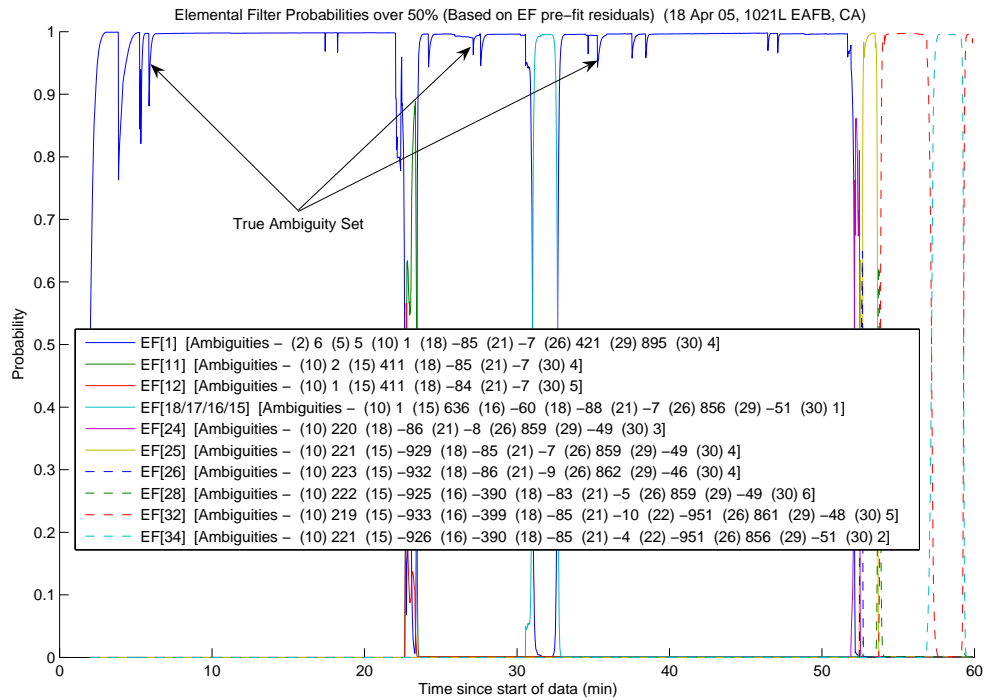


Figure A.83: Case F1.2: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.84 is EF[1], Figure A.85 is EF[21], Figure A.86 is EF[28], Figure A.87 is EF[37], and Figure A.88 is EF[39].

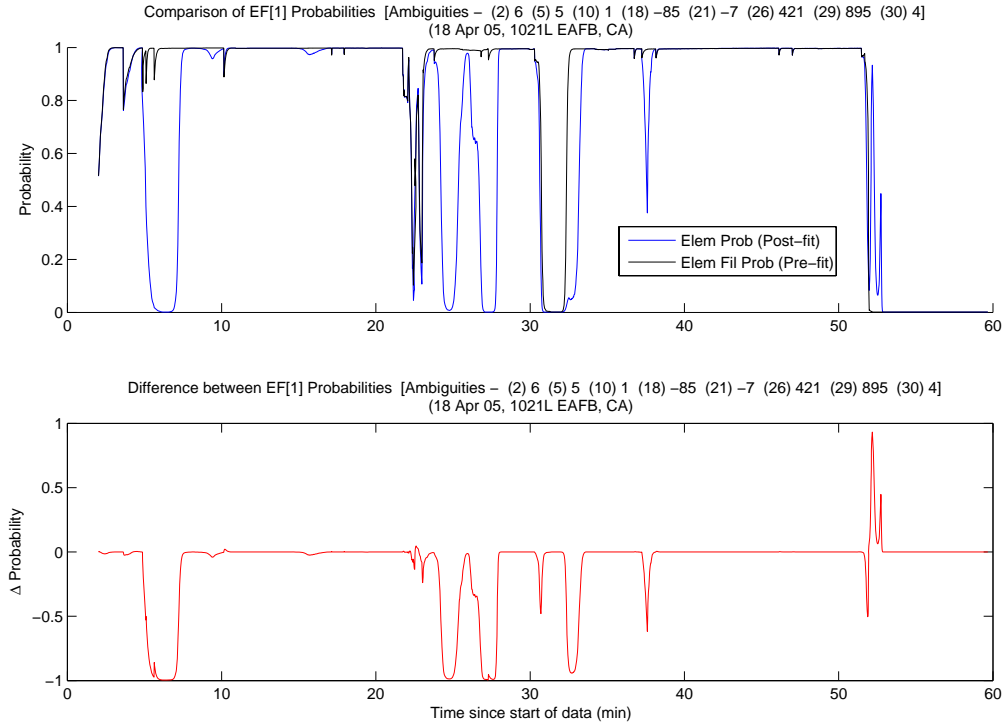


Figure A.84: Case F1.2: EF[1] Probability Comparison

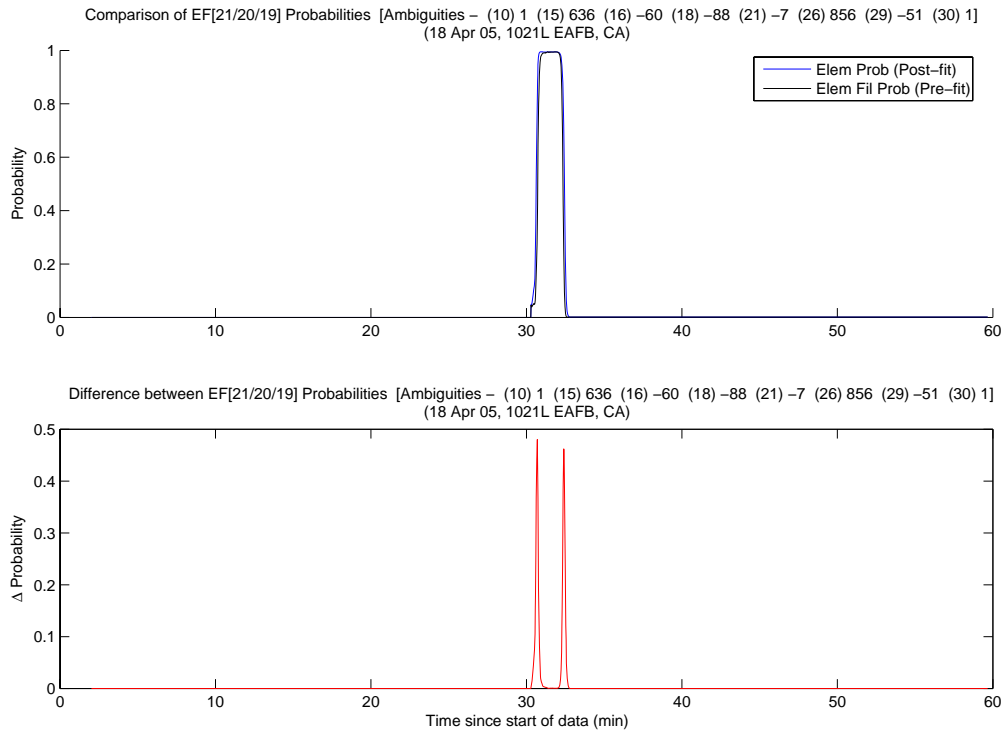


Figure A.85: Case F1.2: EF[21] Probability Comparison

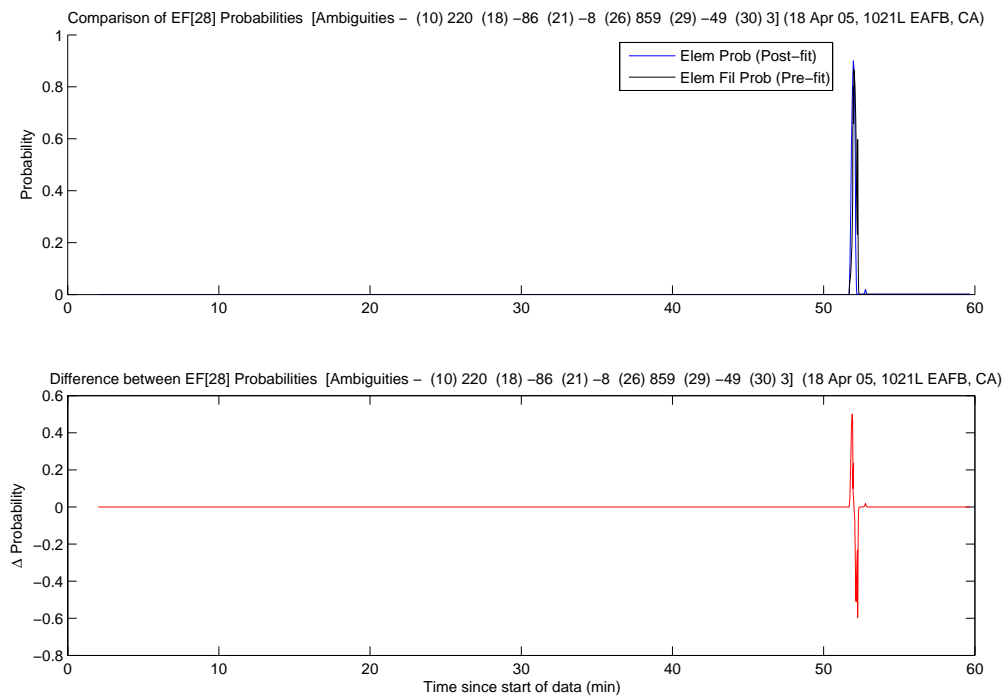


Figure A.86: Case F1.2: EF[28] Probability Comparison

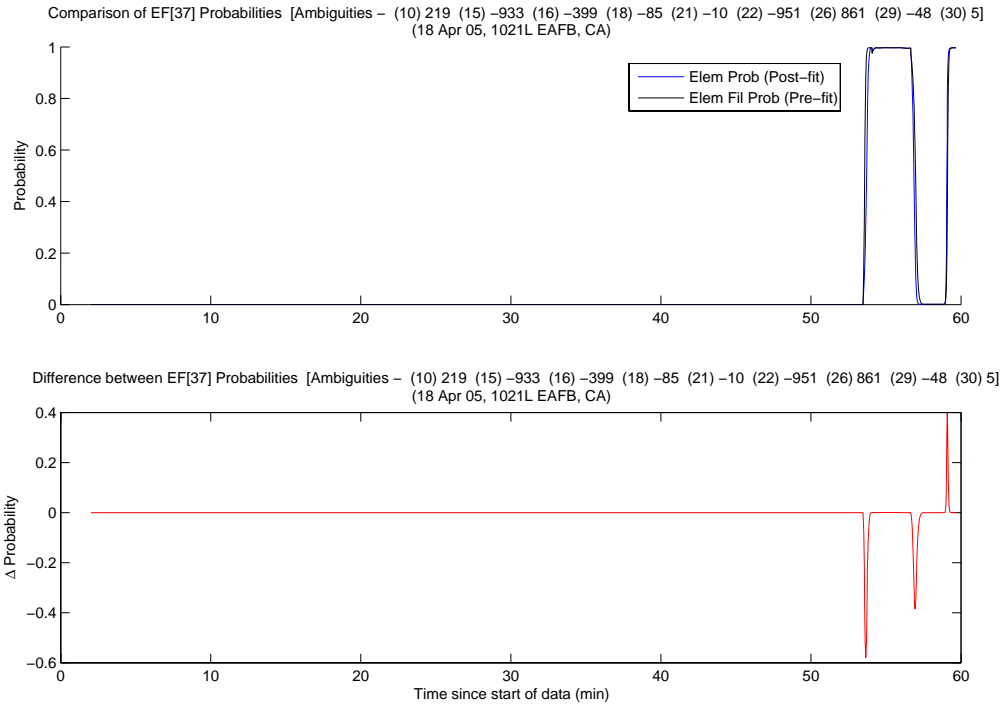


Figure A.87: Case F1.2: EF[37] Probability Comparison

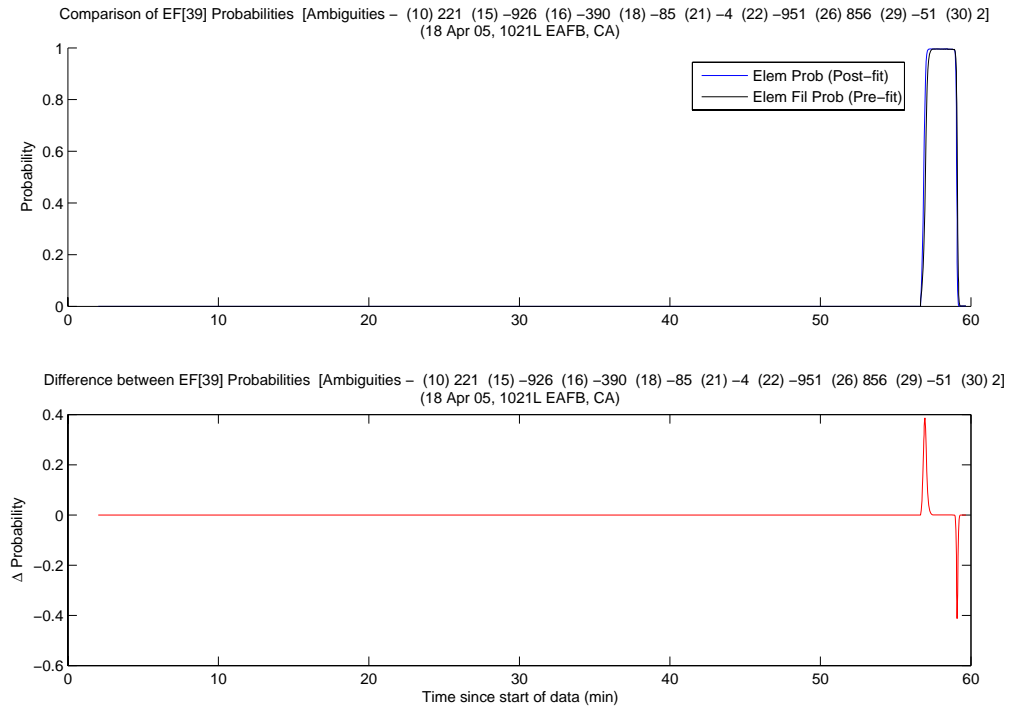


Figure A.88: Case F1.2: EF[39] Probability Comparison

The MMAE position error is shown in Figure A.89. The majority of the error was from the cycle slips at 52 minutes, and at approximately 23 and 30 minutes when erroneous ambiguity set's SOS residuals were lower than the true set.

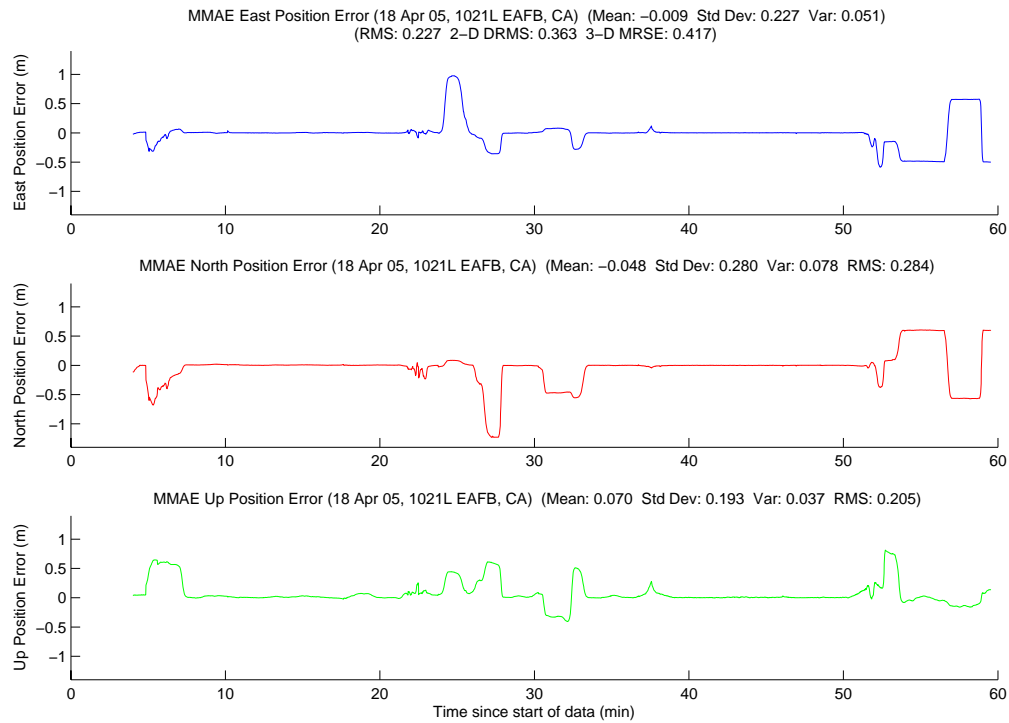


Figure A.89: Case F1.2: MMAE Position Error

Finally, Figure A.90 displays the position error plot for the minimum indicator variable. The impact of the minimum indicator choosing the incorrect ambiguity set at approximately 23 and 52 minutes is evident in the plot. It is interesting to note that, when the minimum indicator chose the wrong set at 23 minutes, the true ambiguity's residuals are well above the incorrectly selected ambiguity's residuals. Also, almost 90 percent of the MMAE probability was absorbed by the same incorrect set. At the same time, the total number of satellites dropped to 5. This was the absolute minimum number of satellites needed to determine the true ambiguity set. Even though the minimum indicator did not have the true set after 52 minutes, the position errors remained reasonable. The errors were similar to having a single ambiguity off by one cycle. Finally, Table A.6 gives a summary of each method's position error.

Even with the incorrect ambiguity set selected during the last eight minutes of the data run, the minimum indicator error remained within the desired centimeter-level accuracy for MRSE.

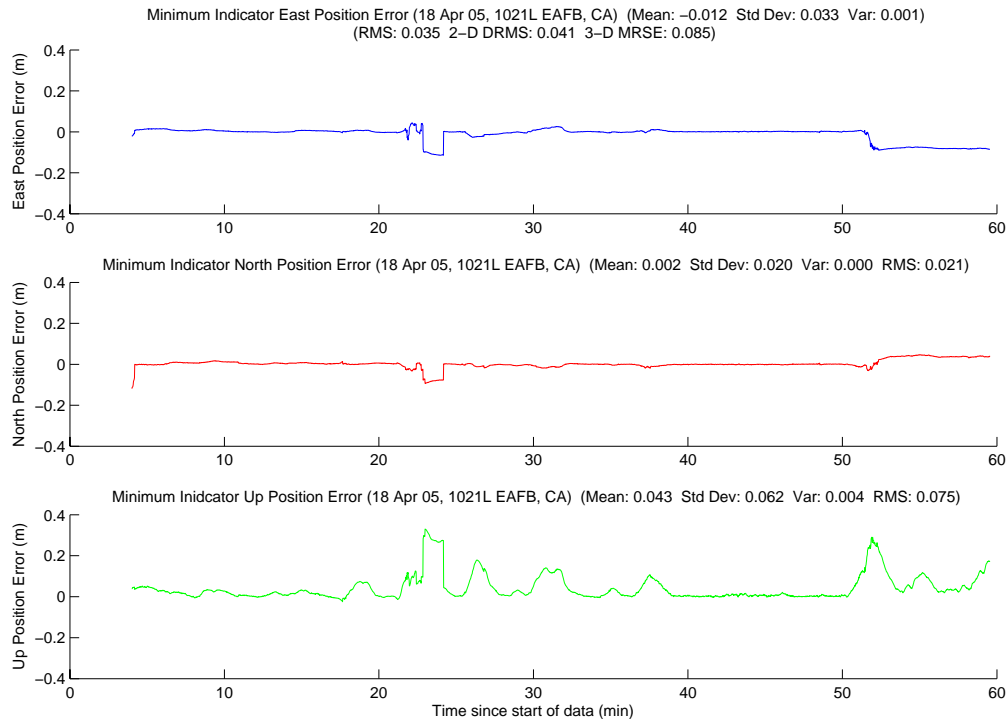


Figure A.90: Case F1.2: Minimum Indicator Position Error

Table A.6: Case F1.2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.008	0.062	-0.117	0.045	0.109	0.193	0.140	0.262
MMAE	-0.009	0.227	-0.048	0.280	0.070	0.193	0.363	0.417
Min. Ind.	-0.012	0.033	0.002	0.020	0.043	0.062	0.041	0.085

A.2.3 Case F2.1: First Data Set Flight Test 2. Case F2.1 was the second of two flight tests at Edwards AFB during April 2005. The flight took off at 0942L. The weather was clear with light winds. Unfortunately, due to Javad hardware problems, only a limited amount of data was collected. Also, the program was terminated prior to the preset runtime, thus not all of the realtime data was saved. The actual data set begins just after engine start, runs through take off, and into the area work. During the test, the aircraft ranged from 19.1 to 1353.9 meters apart. Figure A.91 shows the minimum indicator North, East, Down relative positions.

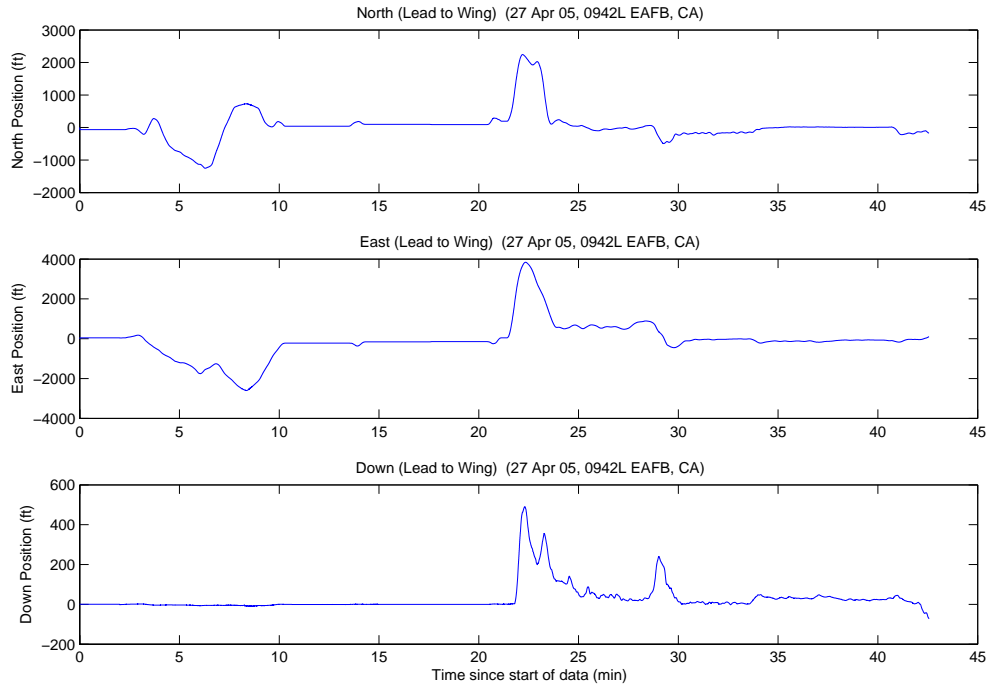


Figure A.91: Case F2.1: North, East, Down Relative Positions

Since not all of the realtime data was saved, information such as the number of unique ambiguity sets returned by the LAMBDA function was unavailable. However, a screen dump was saved for post-flight analysis which included some information. For instance, the minimum indicator correctly chose the true ambiguity set at the first available opportunity (four minutes). Also, the true set was returned from LAMBDA as number one, 2321 of 2415 epochs (96.1 percent).

Figures A.92, A.93, and A.94 show the floating filter position errors. The floating filter reset at 8.5 minutes is due to erroneously detected cycle slips. The aircraft were 810 meters apart while taxiing to the runway when the cycle slips occurred. The data link was having significant problems with numerous lost messages and checksum errors. After enough consecutive messages were lost, the Doppler prediction became less accurate and the difference between adjacent phase measurements was above 1 cycle (which erroneously triggered a cycle slip). The minimum indicator identified the true ambiguity set again, two minutes after the reset at 10.5 minutes. The true set was the same as the first, indicating a cycle slip had not actually occurred.

Take-off was at approximately 22 minutes. This corresponded to an increase in the floating filter's position error and a small bump in the North's estimated 1σ value. Take-off did not seriously affect the floating filter, but it did impact the minimum indicator, which will be discussed later.

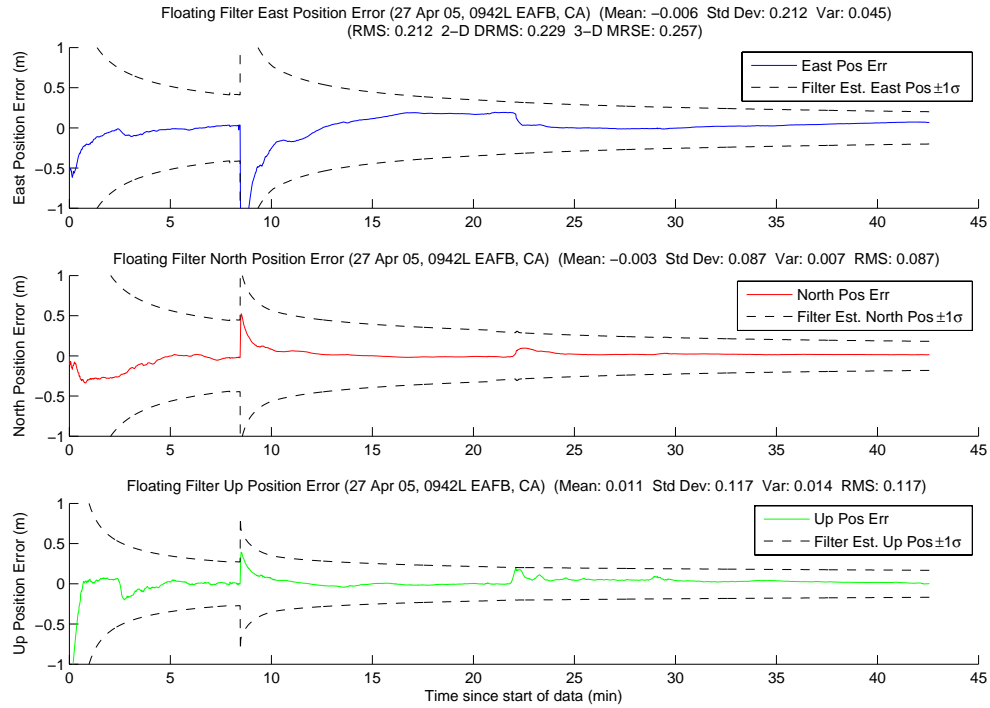


Figure A.92: Case F2.1: Floating Filter Position Errors

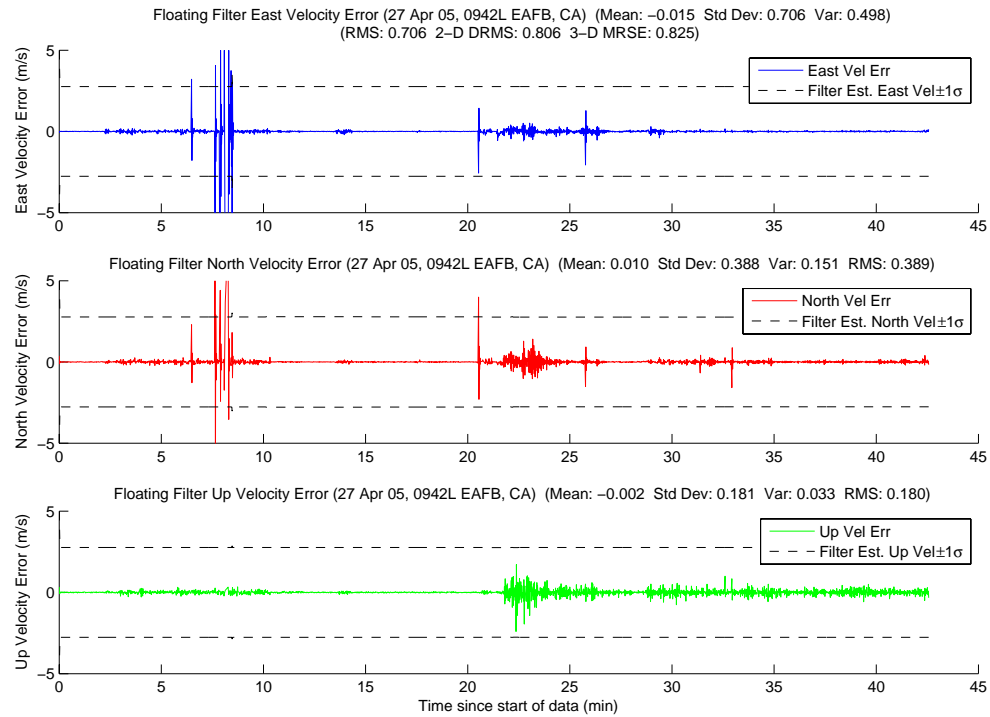


Figure A.93: Case F2.1: Floating Filter Velocity Errors

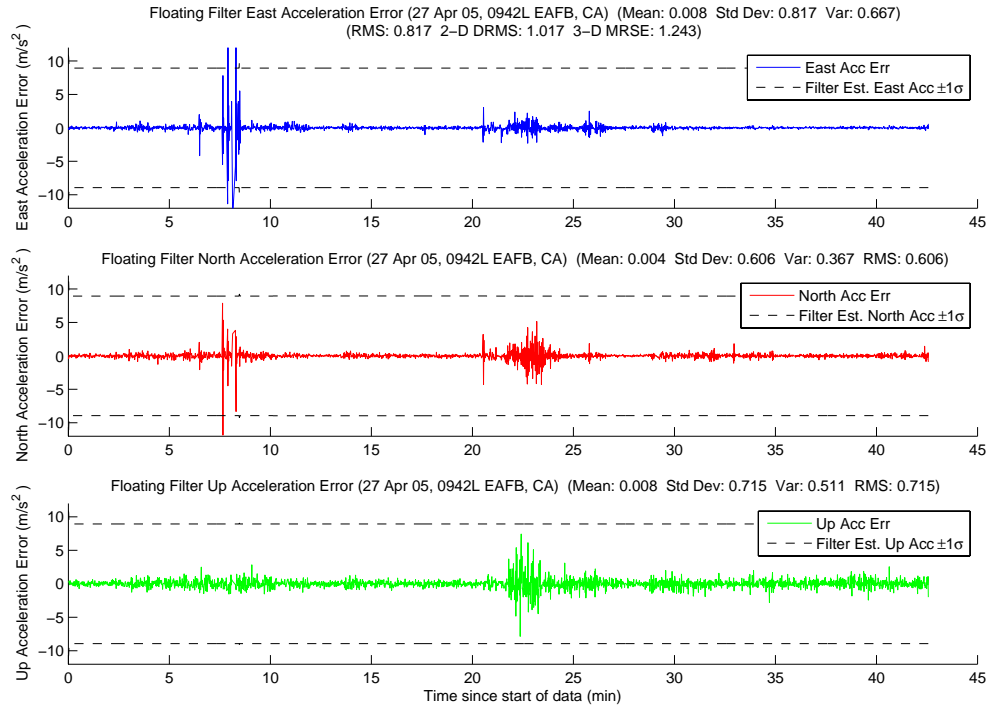


Figure A.94: Case F2.1: Floating Filter Acceleration Errors

Figure A.95 displays PRN 10 measurement information and the floating filter ambiguity estimate. The filter reset is clearly visible in the floating filter ambiguity estimate plot. The phase residuals just prior, during, and just after the reset also indicate how bad the data link connectivity had become. Lastly, it is interesting to note the times when the phase residuals are extremely small. This occurs at the start of the run, at 10.5 minutes, and at 14.5 minutes. As in the previous flight test examples, those times indicate that both aircraft were stopped. The rest of the satellite's plots follow in Figures A.96, A.97, A.98, A.99, A.100, A.101, A.102, A.103, and A.104.

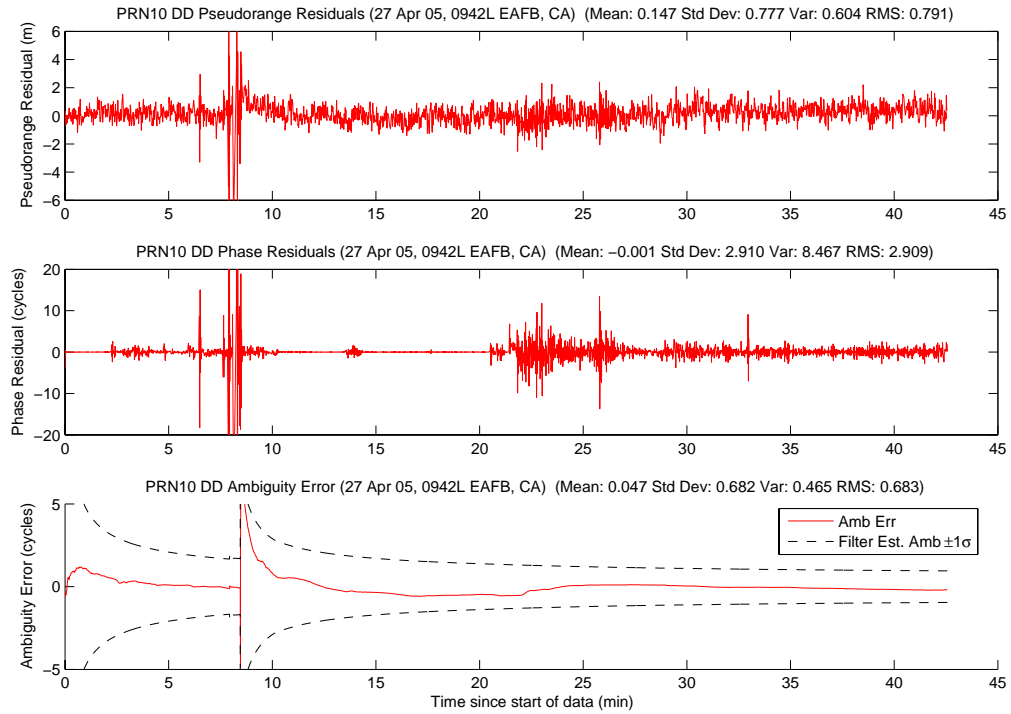


Figure A.95: Case F2.1: Satellite 10 Measurements

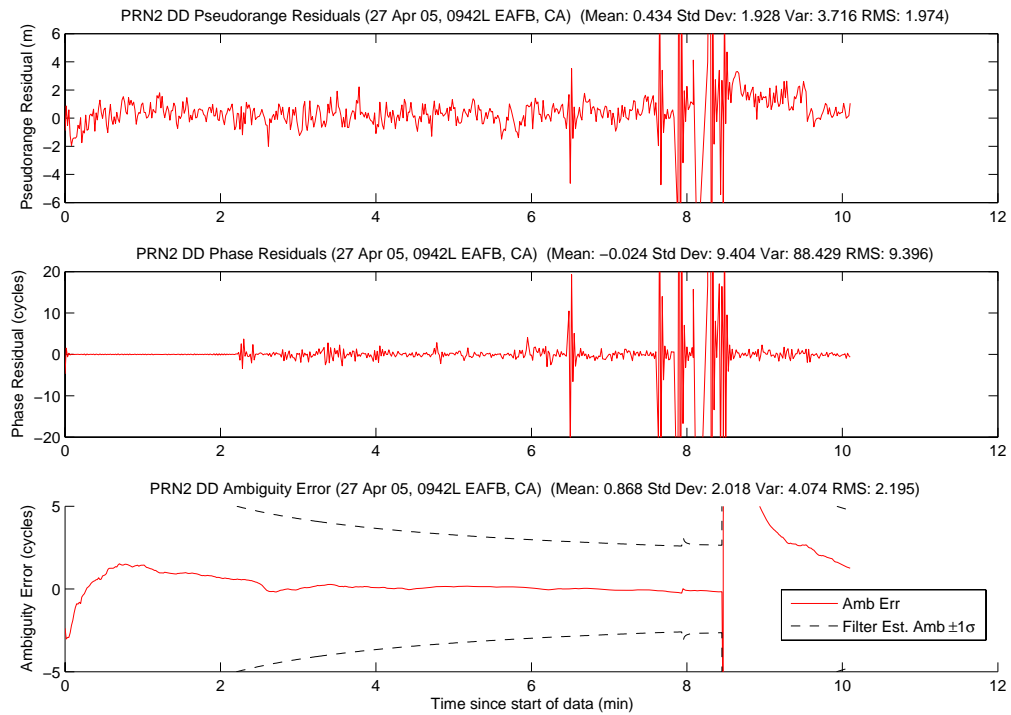


Figure A.96: Case F2.1: Satellite 2 Measurements

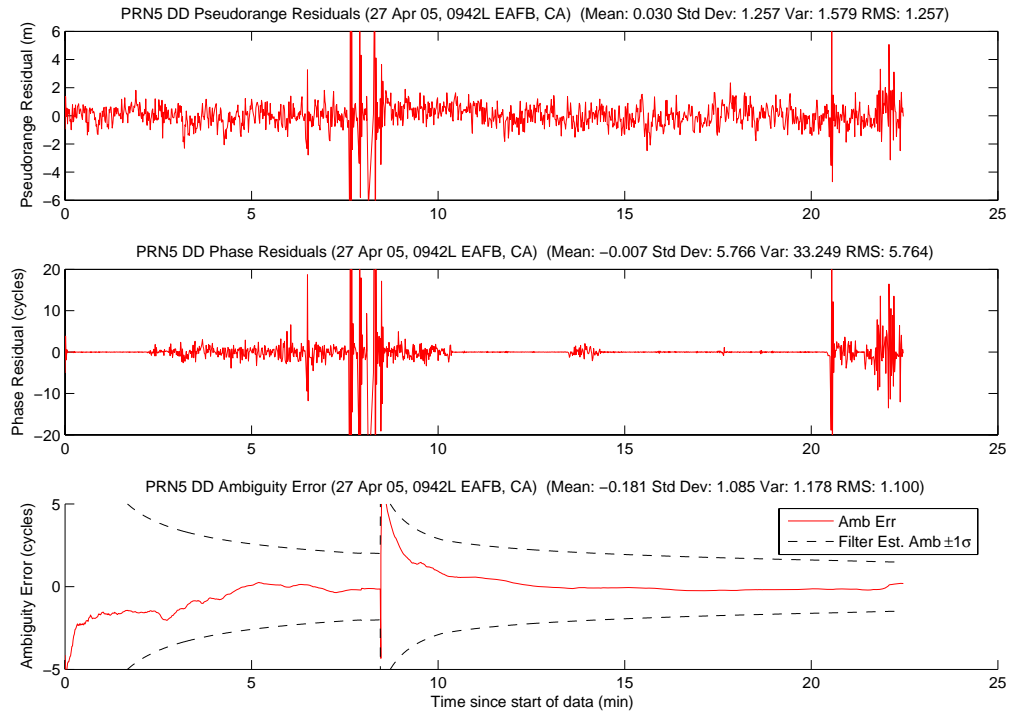


Figure A.97: Case F2.1: Satellite 5 Measurements

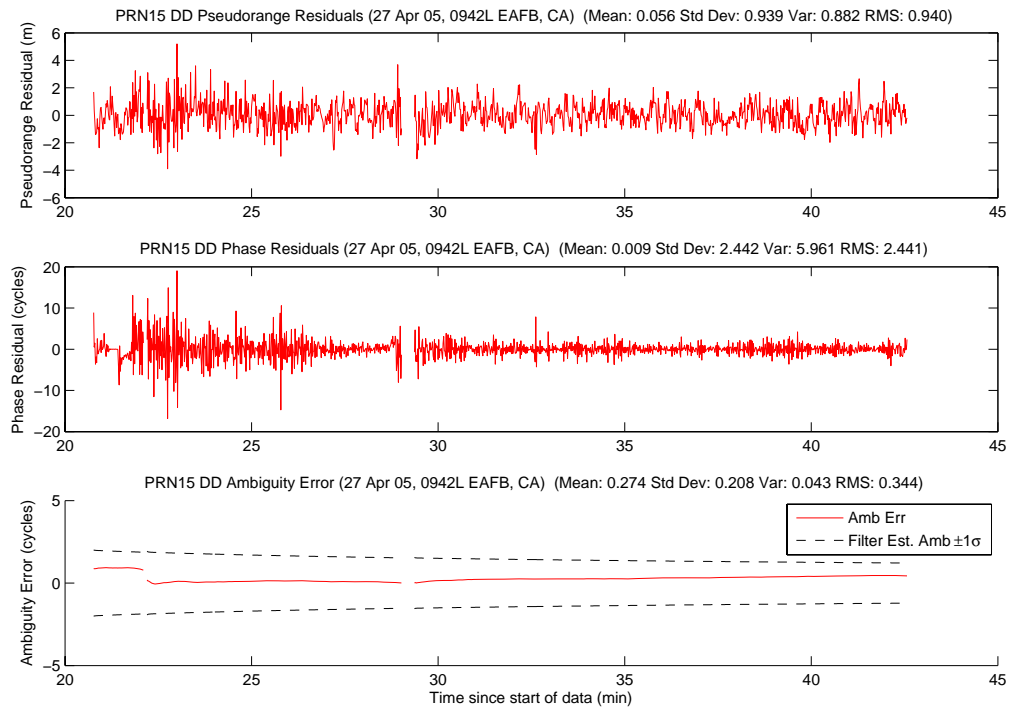


Figure A.98: Case F2.1: Satellite 15 Measurements

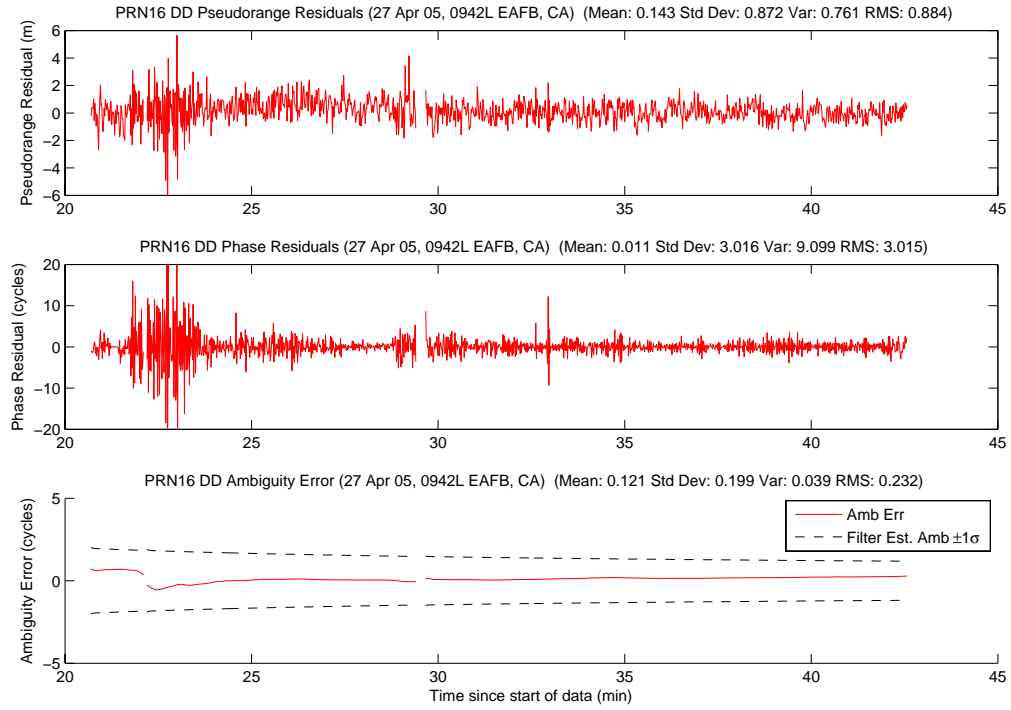


Figure A.99: Case F2.1: Satellite 16 Measurements

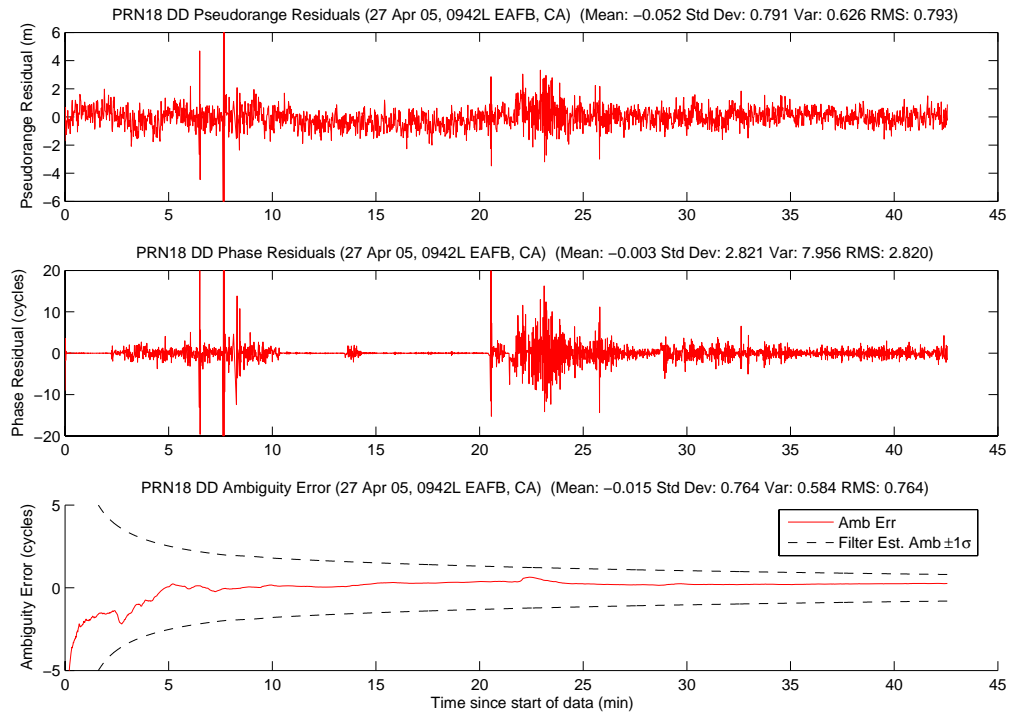


Figure A.100: Case F2.1: Satellite 18 Measurements

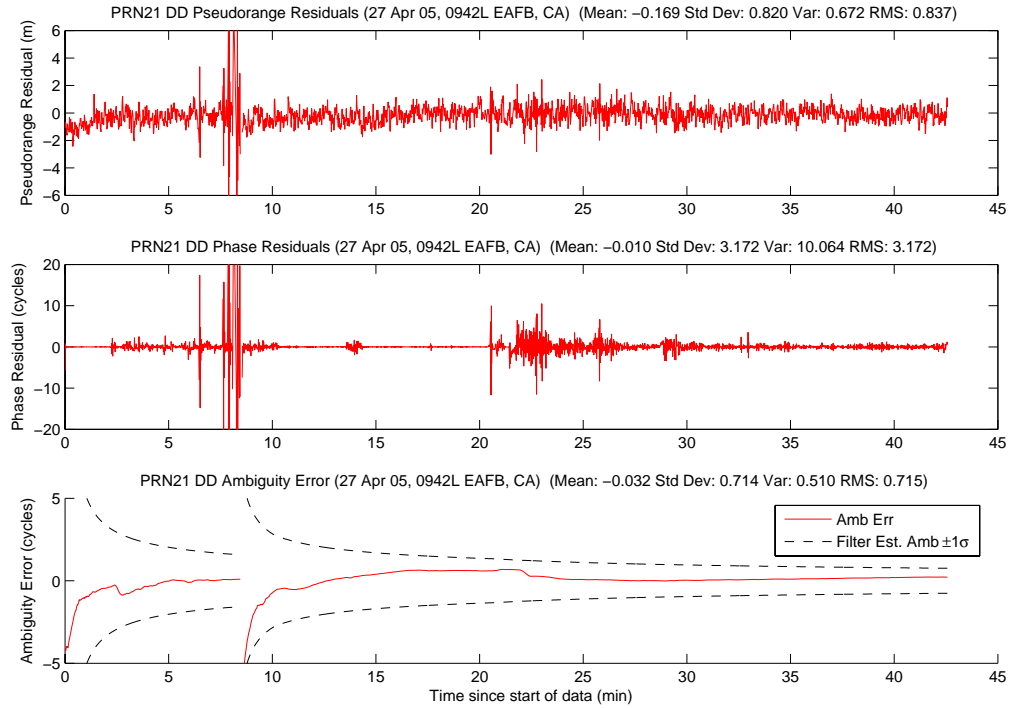


Figure A.101: Case F2.1: Satellite 21 Measurements

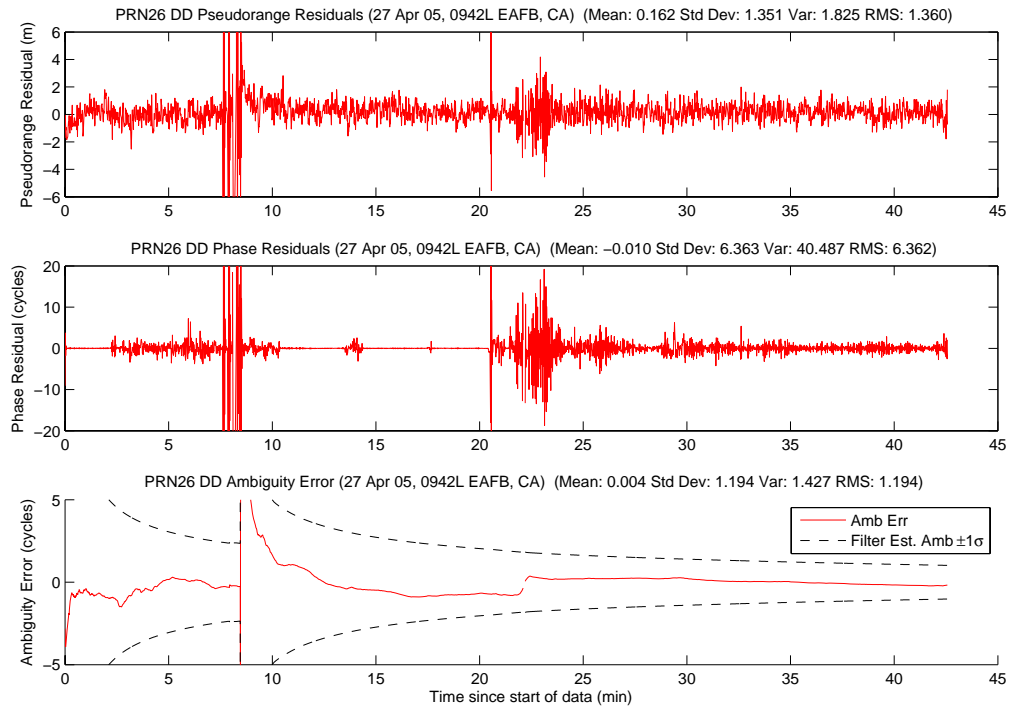


Figure A.102: Case F2.1: Satellite 26 Measurements

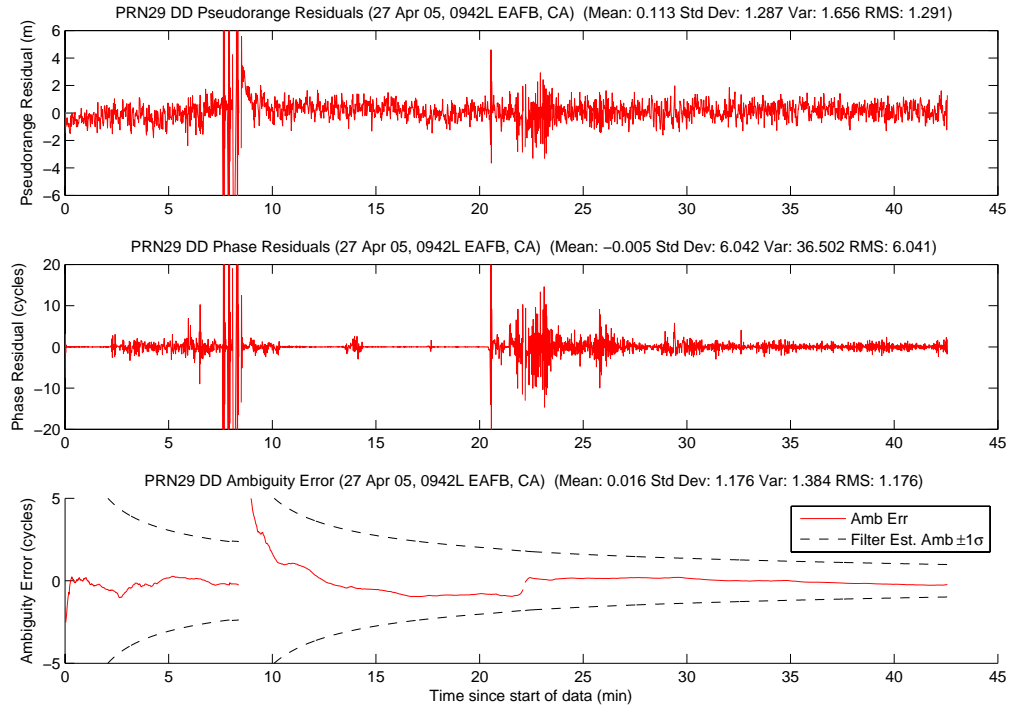


Figure A.103: Case F2.1: Satellite 29 Measurements

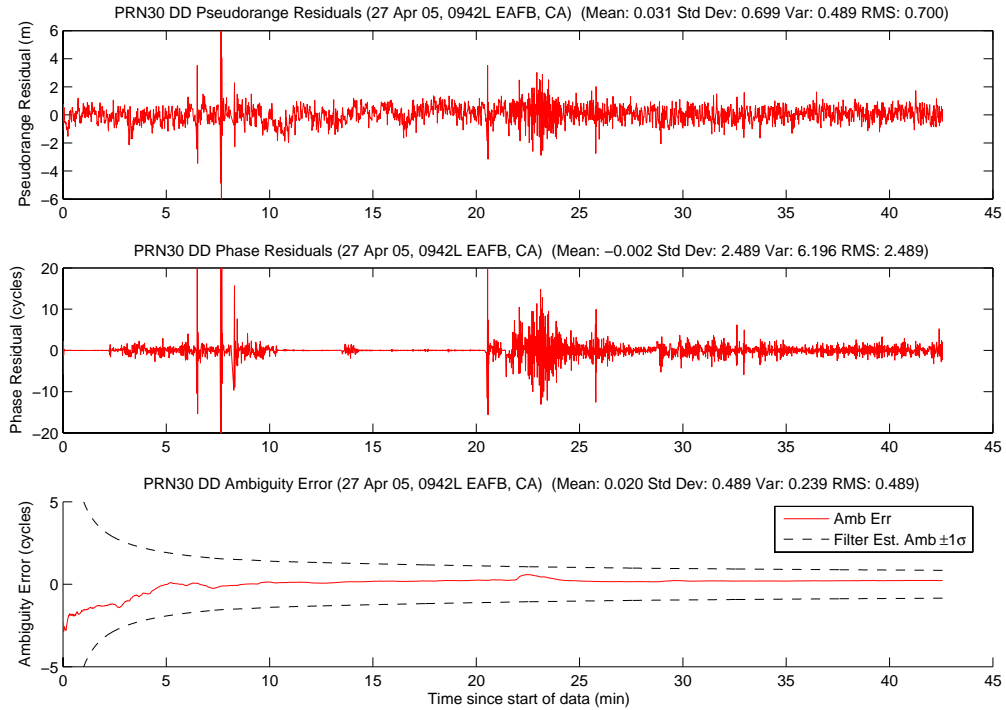


Figure A.104: Case F2.1: Satellite 30 Measurements

Figure A.105 shows the number of visible satellites. The satellite dropouts at 7 minutes were due to erroneous cycle slip detections (the filter reset). The dropouts at 22 minutes, during take-off, were due to actual cycle slips. Finally, the single dropout at 29 minutes was from aircraft maneuvering masking a satellite.

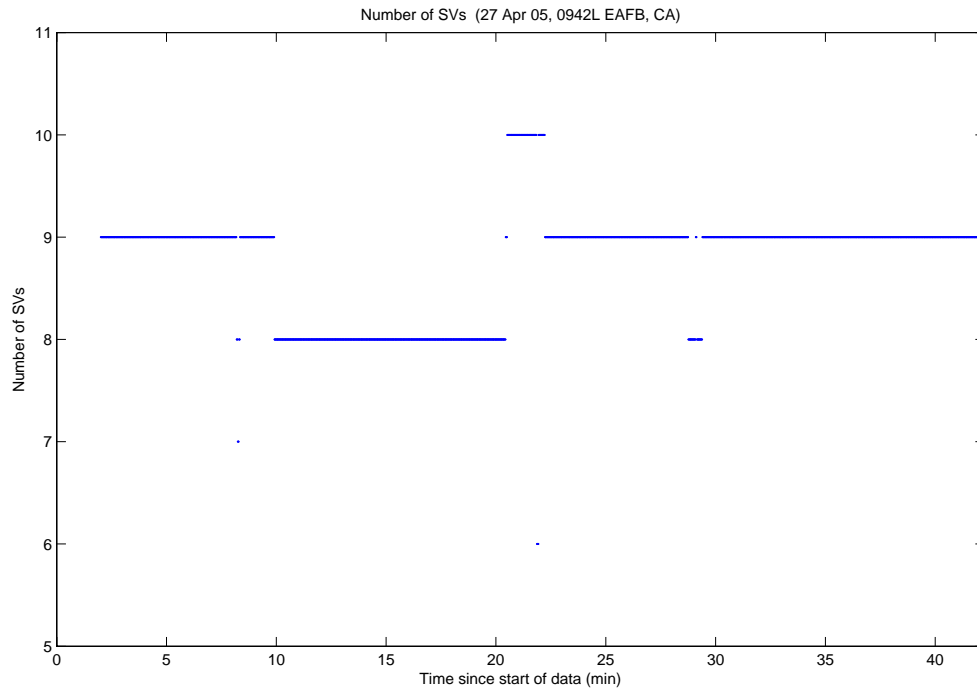


Figure A.105: Case F2.1: Visible Satellites

No SOS residuals plot or probability plot are available since the program was terminated abnormally, but the MMAE position error is shown in Figure A.106. The majority of the error occurred at 8 minutes when the filter was reset, and at 22 minutes during take-off.

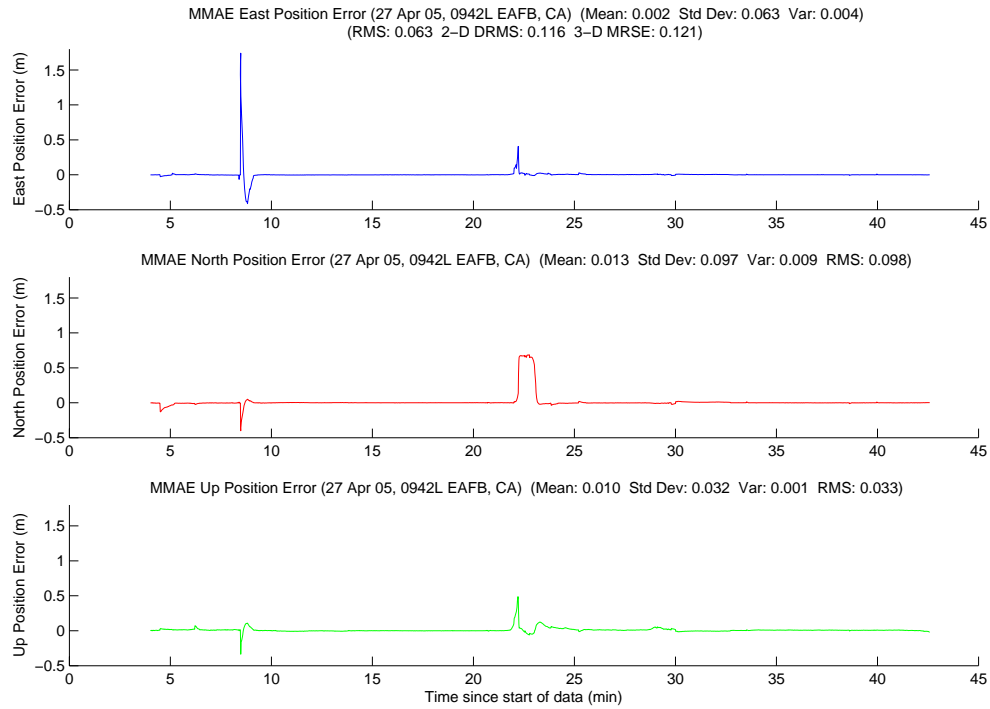


Figure A.106: Case F2.1: MMAE Position Error

Finally, Figure A.107 displays the position error plot for the minimum indicator variable. The minimum indicator correctly chose the true ambiguity set at the beginning of the data run and after the floating filter reset. However, during take-off, four satellites experienced cycle slips. This reduced the number of usable satellites to six. The minimum indicator changed to an incorrect set shortly after take-off, at 22.167 minutes. This is the first big jump in the minimum indicator position error. The minimum indicator quickly changed to another erroneous ambiguity set 30 seconds later (the second jump in the position error). It remained on this set for 6 minutes, until it found the true ambiguity set once more. Unfortunately, the minimum indicator changed again to a set that was off by one cycle on PRN 15 a minute later. Ten seconds later, the minimum indicator transitioned to another incorrect set (it was off by 1 cycle on three PRNs). Then, 30 seconds later it found the true ambiguity set again. The minimum indicator did not change after it found the true set for the last time. Although, the minimum indicator incorrectly chose four wrong ambiguity sets, its MRSE was still well within the centimeter-level requirement and half as large as the MMAE. Finally, Table A.7 gives a summary of each method's position error.

Table A.7: Case F2.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.006	0.212	-0.003	0.087	0.011	0.117	0.229	0.257
MMAE	0.002	0.063	0.013	0.097	0.010	0.032	0.116	0.121
Min. Ind.	-0.005	0.037	0.009	0.021	-0.011	0.042	0.044	0.062

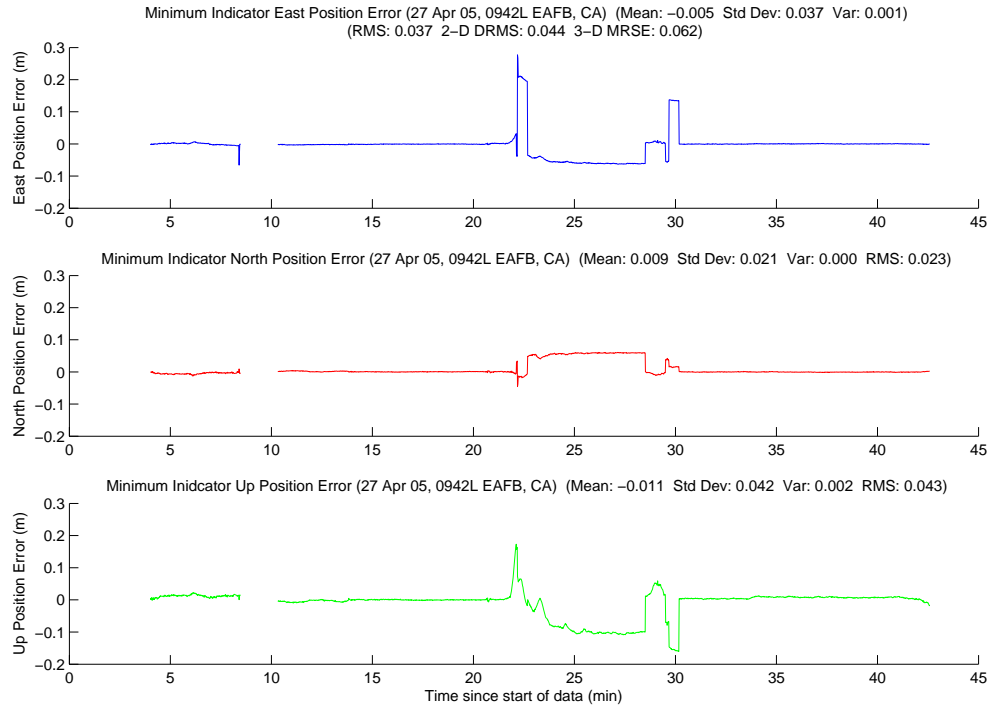


Figure A.107: Case F2.1: Minimum Indicator Position Error

A.2.4 Case F2.2: Flight Test 2, Second Data Set. Case F2.2 started 30 minutes after case F2.1. There were numerous attempts to get the hardware to function properly before this run began; however, it only worked for about 24 minutes on this run. The data run began at 1103L and the aircraft ranged from 155.2 to 1591.6 meters apart. Figure A.108 shows the minimum indicator North, East, Down relative positions.

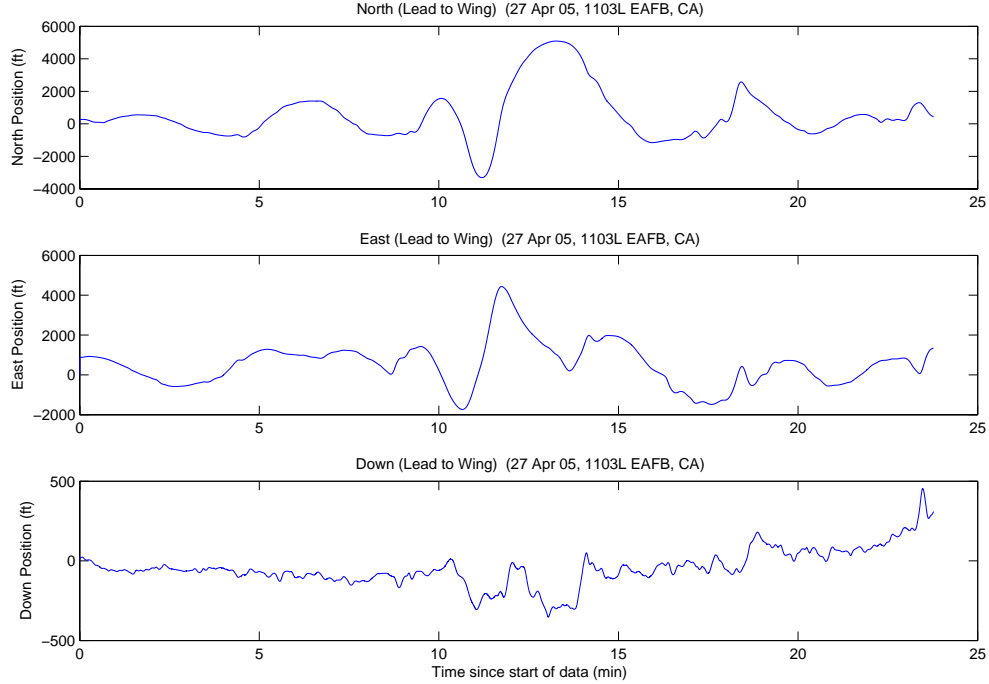


Figure A.108: Case F2.2: North, East, Down Relative Positions

As with case F2.1, the program was terminated prior to the specified runtime, and not all of the realtime data was saved. However, the post-processing data was almost identical to the realtime data and will be displayed where specified.

Post-processing data was used to collect the following information. The LAMBDA function returned a total of 21 unique ambiguity sets during the data run. The correct set was the top candidate returned from the first LAMBDA pair. The minimum indicator correctly chose the true ambiguity set at 4 minutes. It did not change throughout the data run. The true set was returned from LAMBDA as number one, 1327 of 1327 epochs (100.0 percent), with a maximum ratio of 109.8.

The rest of the plots are based on realtime data, except the SOS residuals and MMAE probability plots. Figures A.109, A.110, and A.111 show the floating filter errors. The floating filter performed as expected based on initializing with nine satellites.

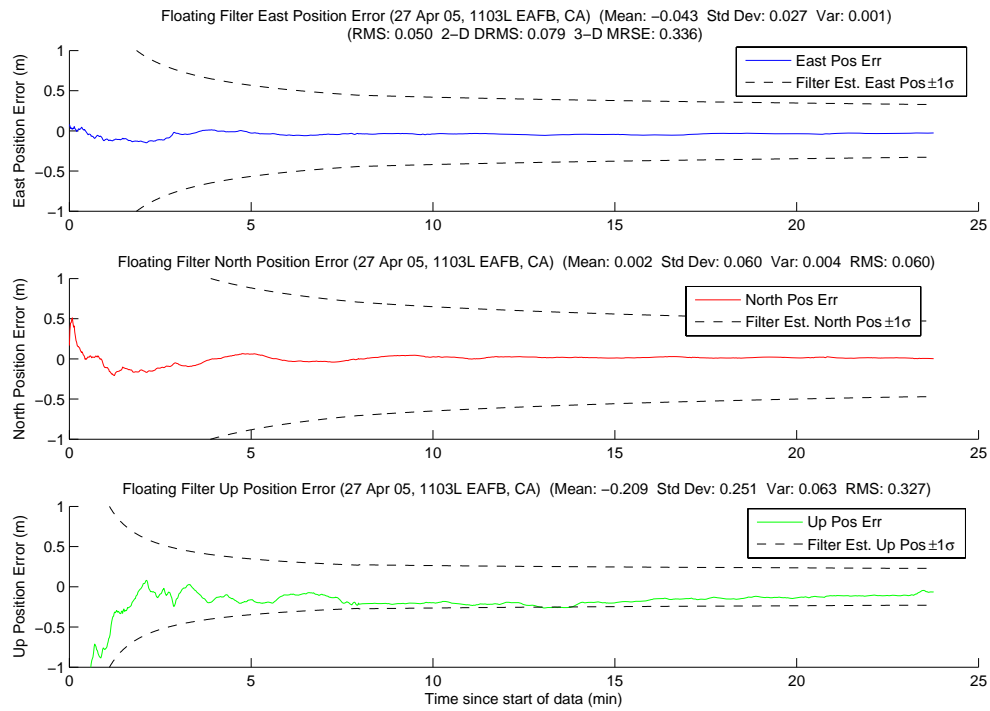


Figure A.109: Case F2.2: Floating Filter Position Errors

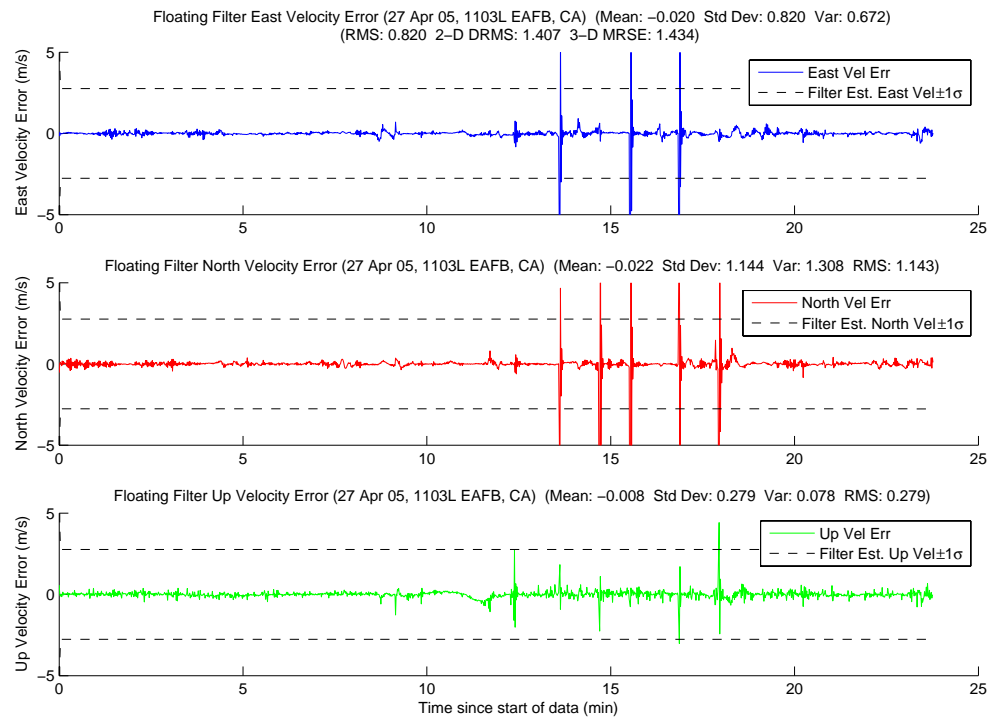


Figure A.110: Case F2.2: Floating Filter Velocity Errors

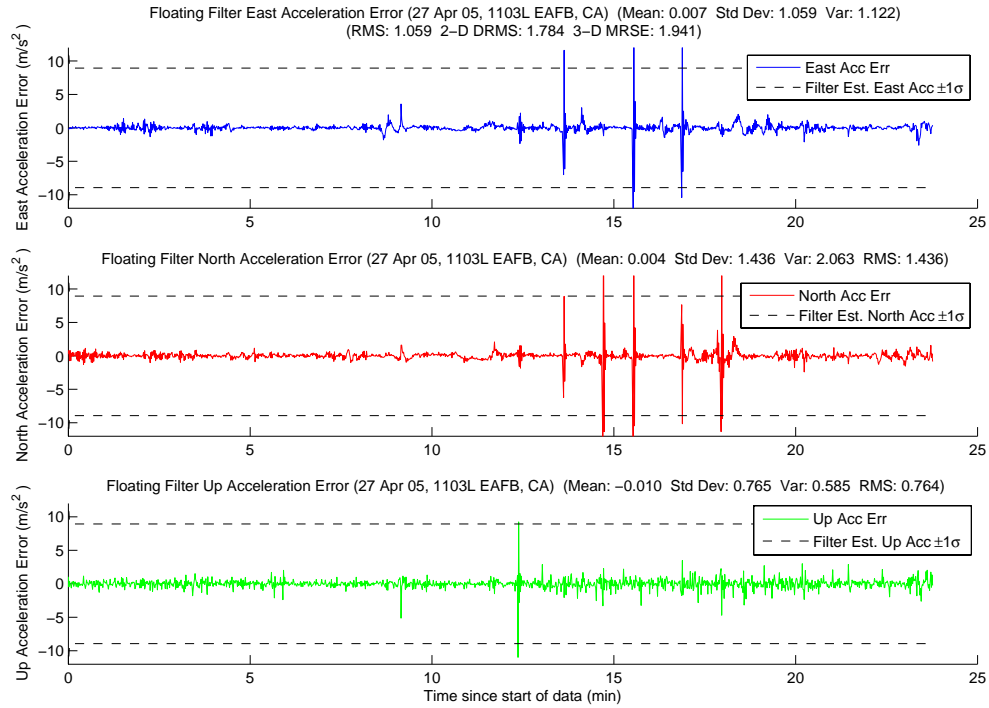


Figure A.111: Case F2.2: Floating Filter Acceleration Errors

Figure A.112 displays PRN 15 measurement information and the floating filter ambiguity estimate. The decrease in the floating filter 1σ ambiguity estimate at approximately 12 minutes was due to a change in the maximum elevated satellite. The spikes in the code / phase residuals were due to aircraft maneuvering. The rest of the satellite's plots follow in Figures A.113, A.114, A.115, A.116, A.117, A.118, A.119, and A.120.

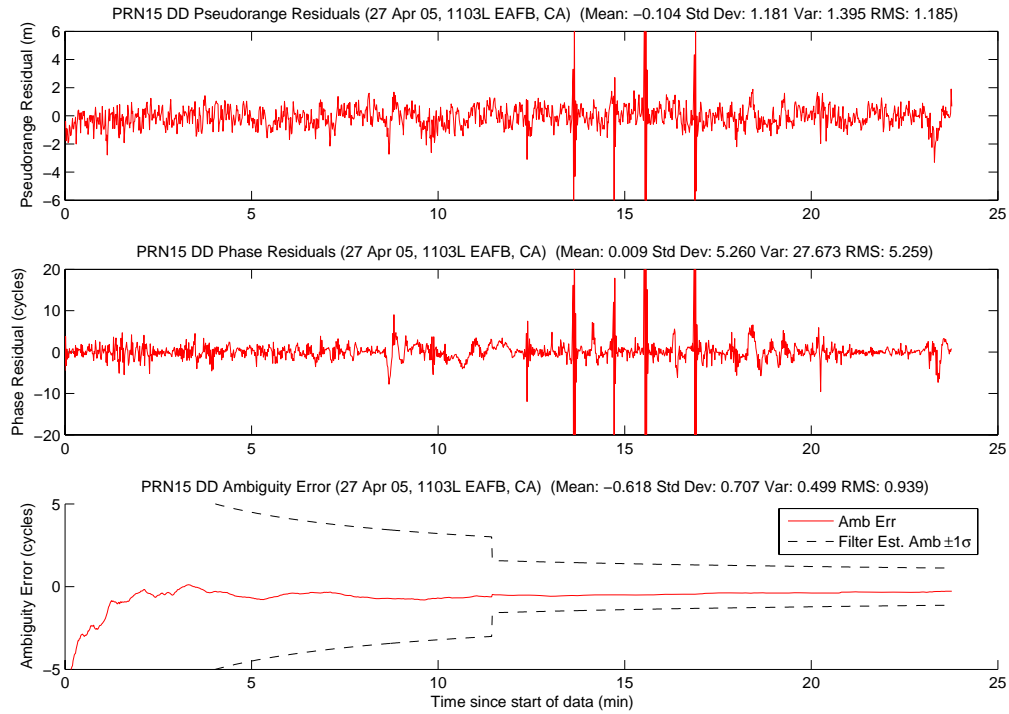


Figure A.112: Case F2.2: Satellite 15 Measurements

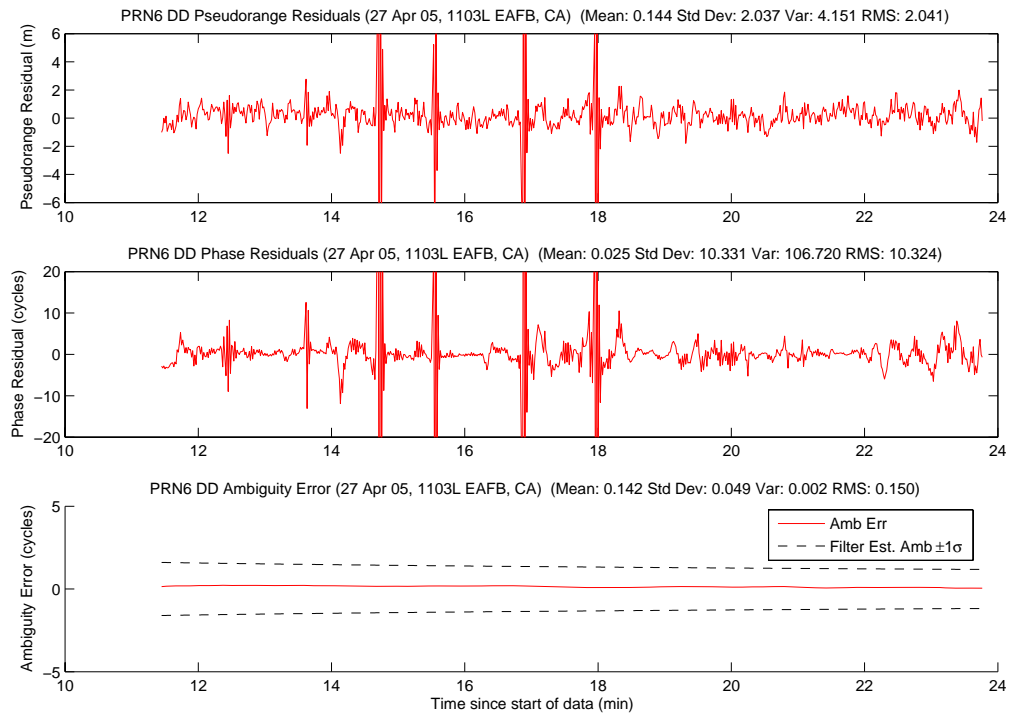


Figure A.113: Case F2.2: Satellite 6 Measurements

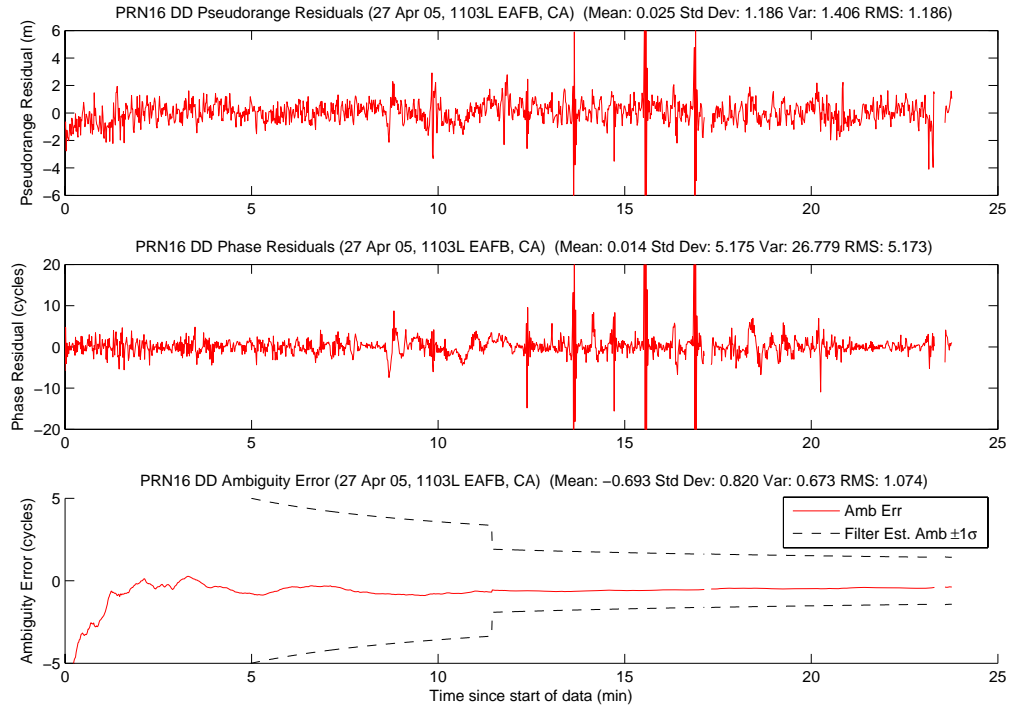


Figure A.114: Case F2.2: Satellite 16 Measurements

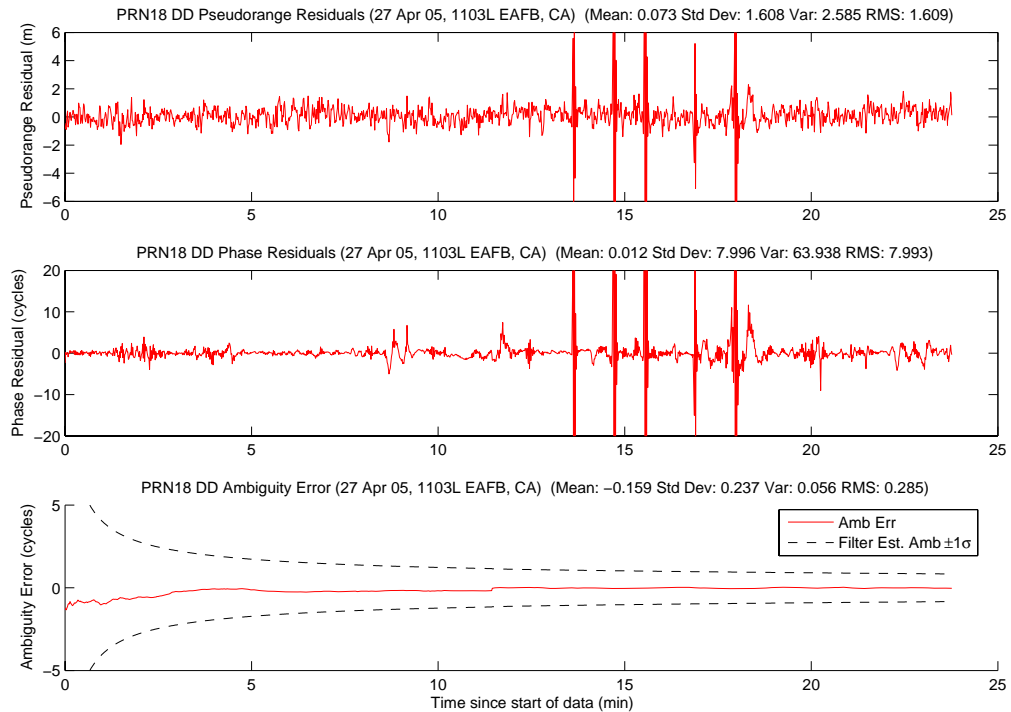


Figure A.115: Case F2.2: Satellite 18 Measurements

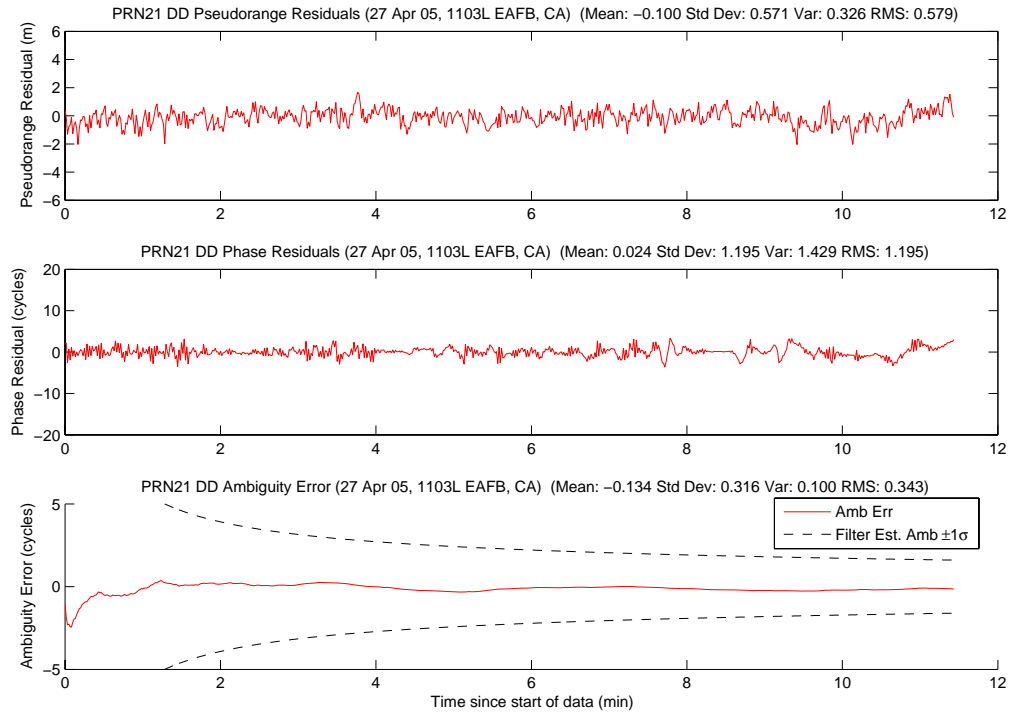


Figure A.116: Case F2.2: Satellite 21 Measurements

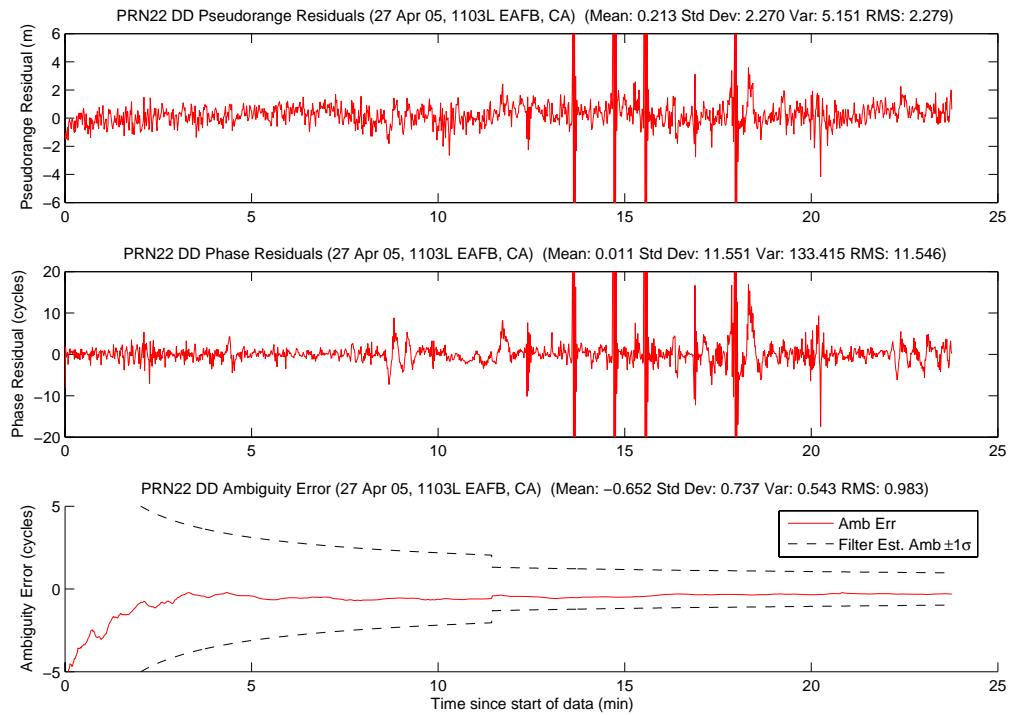


Figure A.117: Case F2.2: Satellite 22 Measurements

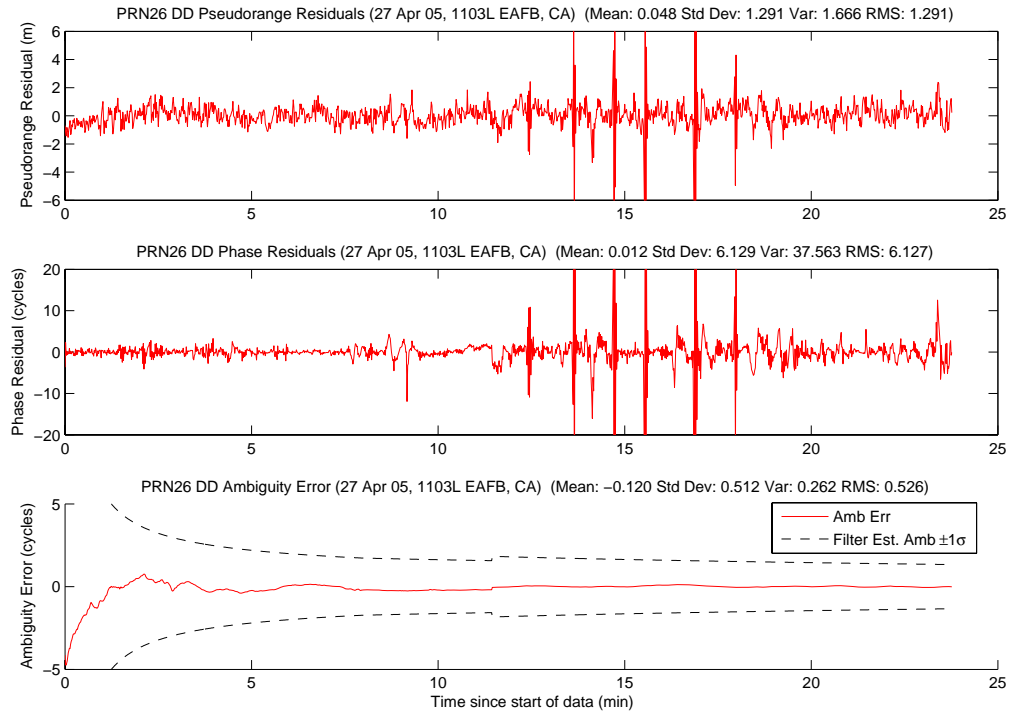


Figure A.118: Case F2.2: Satellite 26 Measurements

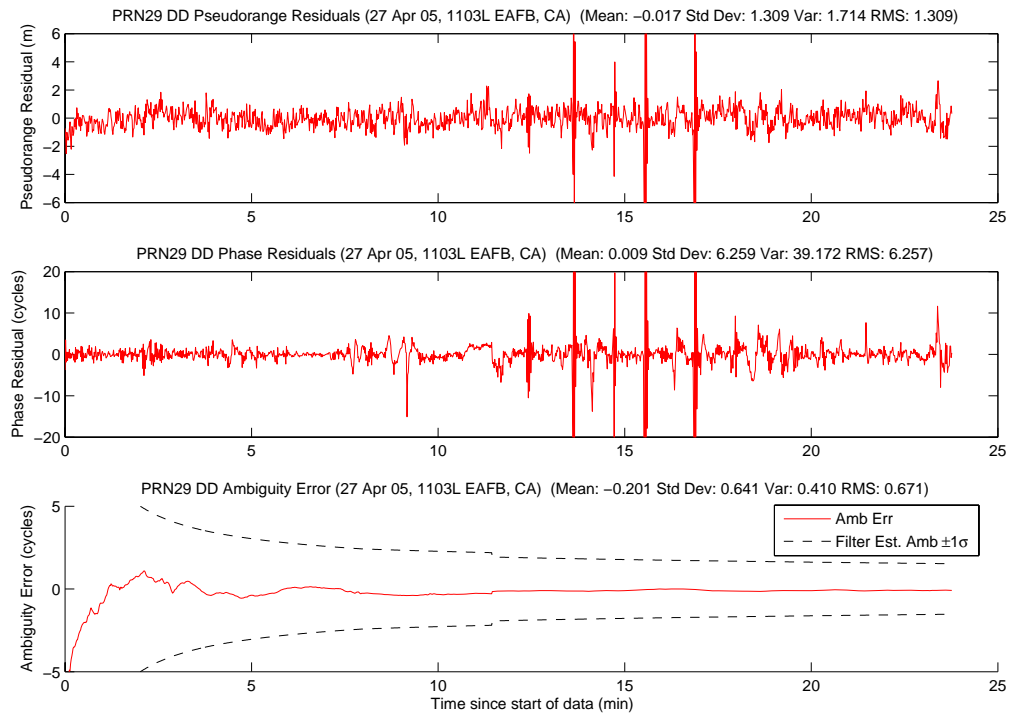


Figure A.119: Case F2.2: Satellite 29 Measurements

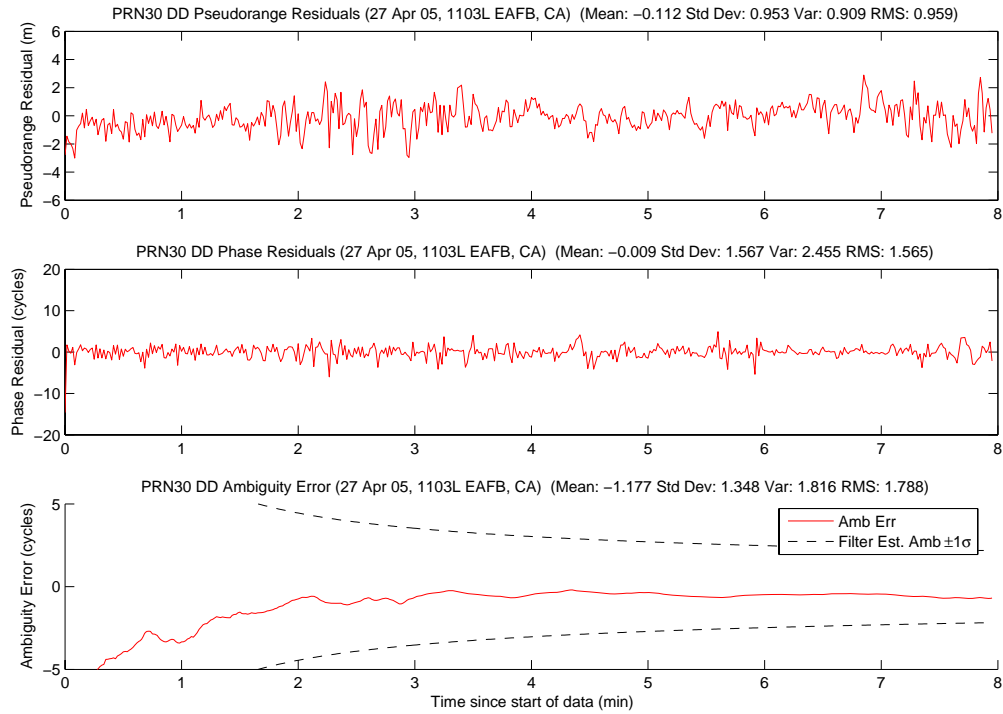


Figure A.120: Case F2.2: Satellite 30 Measurements

Figure A.121 shows the number of visible satellites. The satellite dropouts at 17 and 23 minutes were due to aircraft maneuvering masking a satellite (PRN 16).

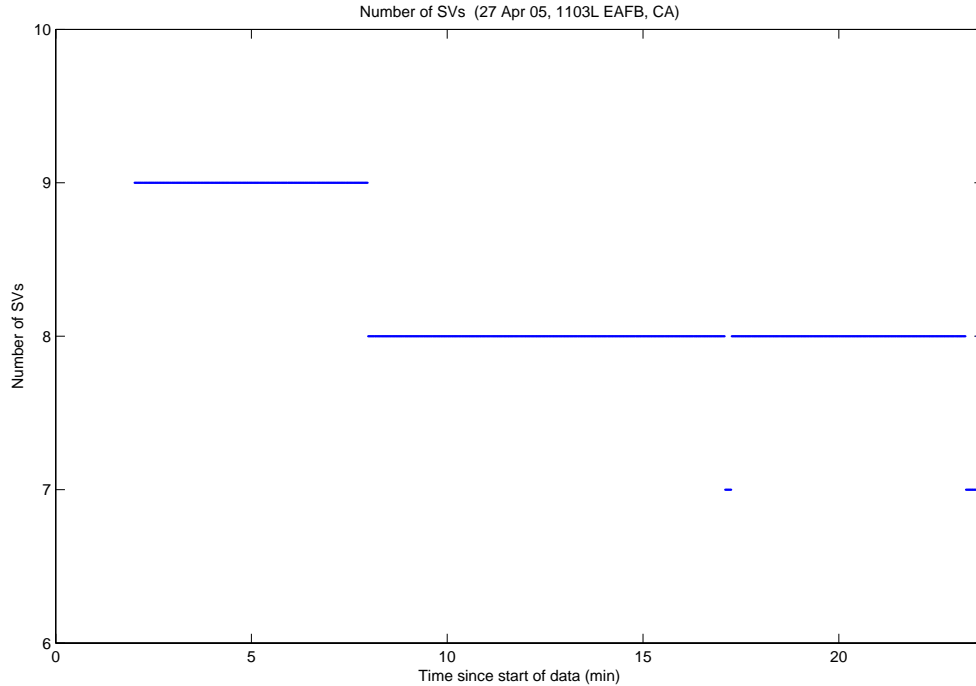


Figure A.121: Case F2.2: Visible Satellites

The SOS residuals plot and probability plot were constructed from post-processed data. Figure A.122 shows all of the candidate ambiguity set's SOS residuals. The increase in the true set's SOS residuals at 11.5 minutes occurred when the aircraft were maneuvering 1.4 km apart. Figure A.123 displays the MMAE conditional probability. The true ambiguity set maintained over 49 percent of the total probability until 13 minutes. Then, an erroneous set slowly absorbed 99 percent of the probability. Once the wrong ambiguity set reached 99 percent at 15 minutes, the true set slowly began reabsorbing the probability back. An examination of the SOS residual plot shows why this phenomenon occurred. The erroneous set's SOS residuals are highlighted by gold open squares in Figure A.122. The erroneous set's residuals remain lower, on a parallel course, to the true set's residuals starting at approximately 13 minutes. Once the wrong set's residuals start growing again at 15 minutes, and the true set levels off, the true ambiguity set begins to reabsorb the probability. This was a classic example of how an incorrect ambiguity set can have a local minimum below the true ambiguity set.

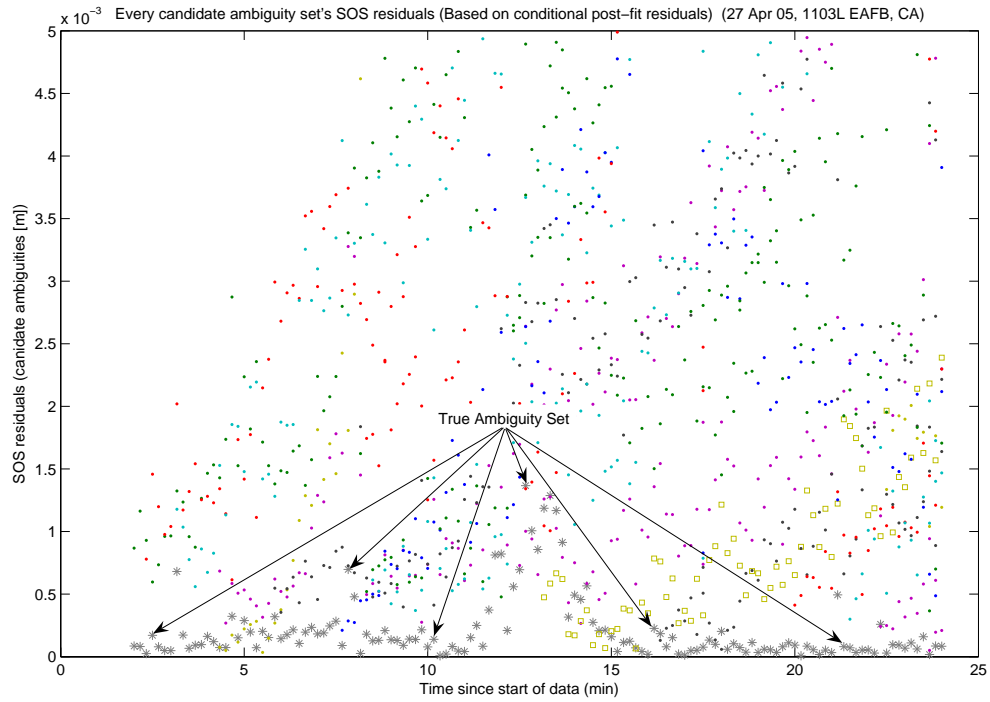


Figure A.122: Case F2.2: SOS Residuals

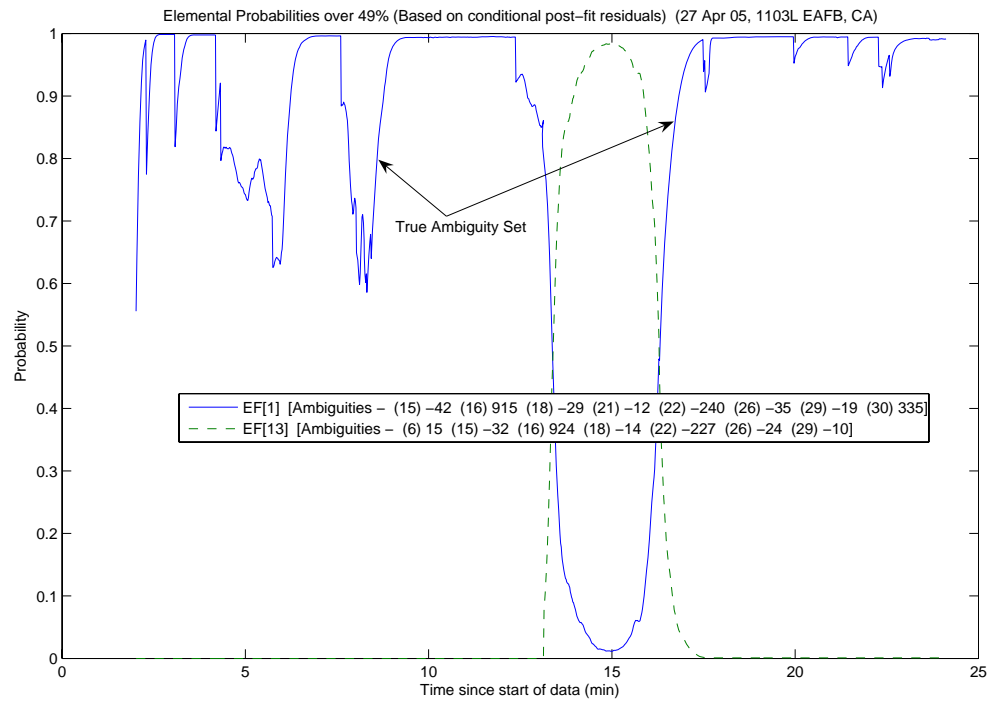


Figure A.123: Case F2.2: Conditional Probabilities for Selected MMAE Elemental Filters

The following figure shows the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.124 is EF[1].

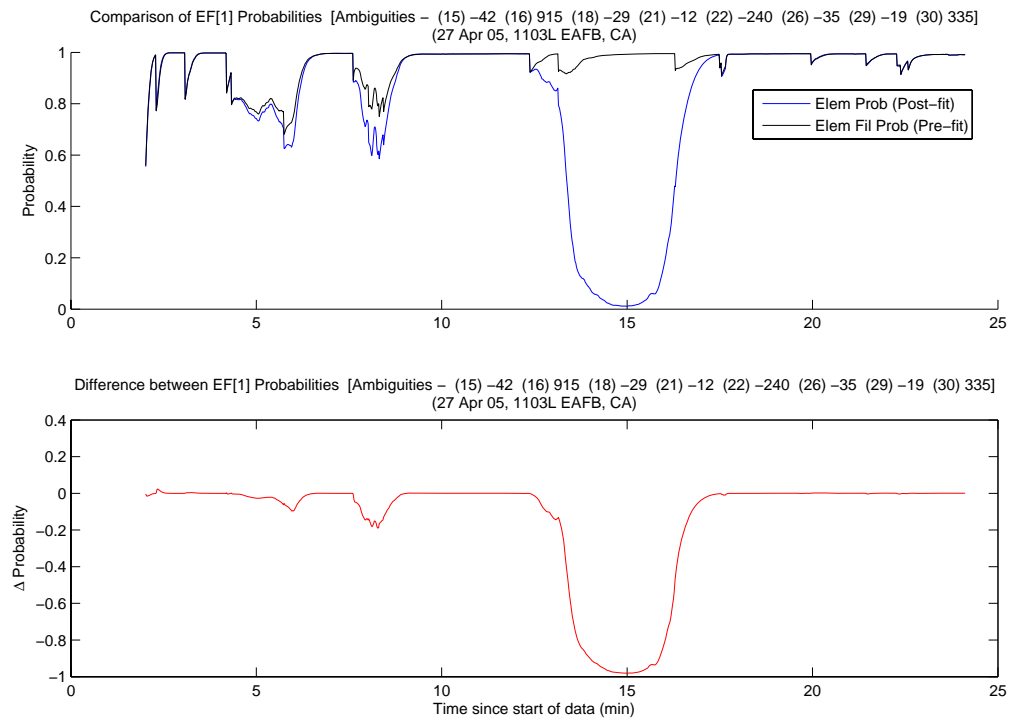


Figure A.124: Case F2.2: EF[1] Probability Comparison

The MMAE position error is shown in Figure A.125. The probability shift from above is visible in the growth of the MMAE position error from 12 to 16 minutes. The total error grew to almost 2 meters during the four minute period, an unacceptable amount for precise formation flying. Also, this was the major draw back to the MMAE blended solution and the reason the minimum indicator method was developed.

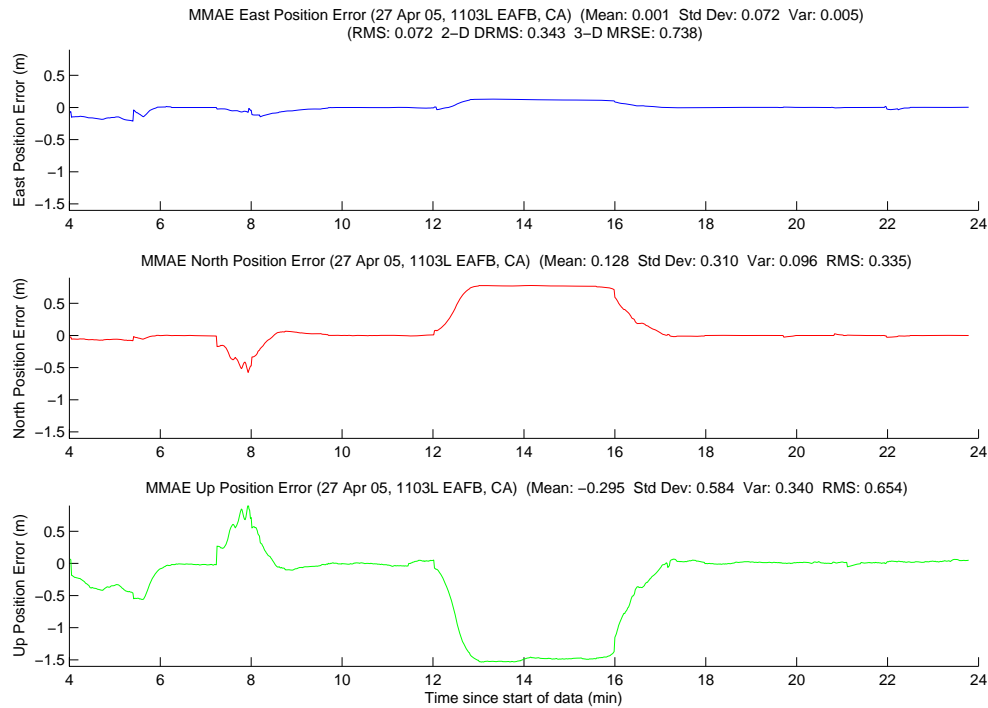


Figure A.125: Case F2.2: MMAE Position Error

Finally, Figure A.126 displays the position error plot for the minimum indicator variable. The minimum indicator MRSE remained at 2.2 cm for the 24 minute run, well within the centimeter-level requirement. The minimum indicator was always “locked” to the true ambiguity set. Lastly, Table A.8 gives a summary of each method’s position error.

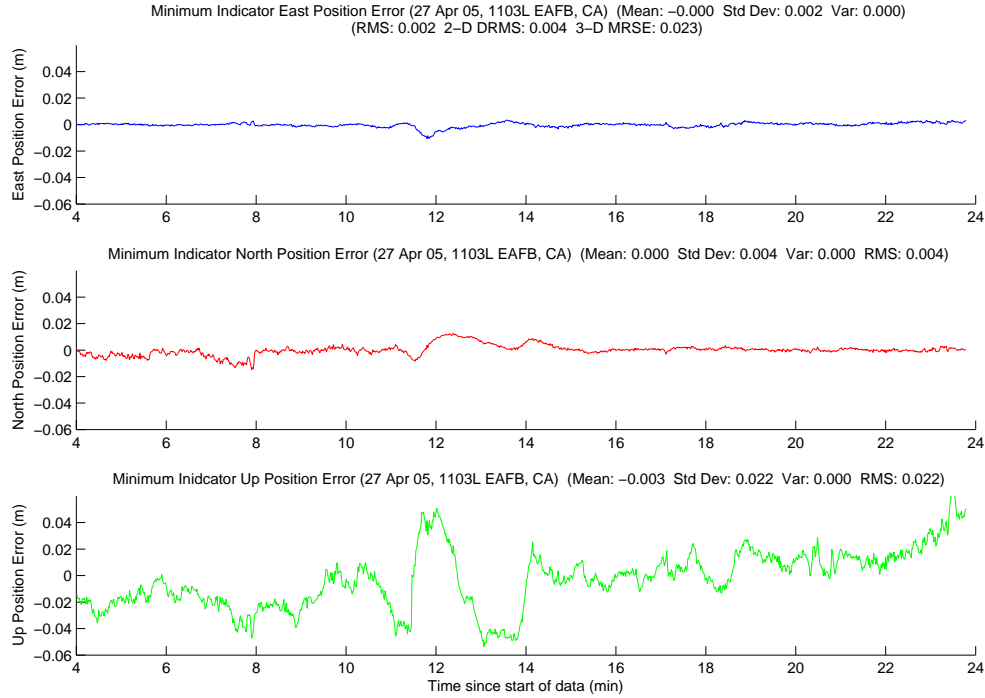


Figure A.126: Case F2.2: Minimum Indicator Position Error

Table A.8: Case F2.2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.043	0.027	0.002	0.060	-0.209	0.251	0.079	0.336
MMAE	-0.001	0.072	0.128	0.310	-0.295	0.584	0.343	0.738
Min. Ind.	-0.000	0.002	0.000	0.004	-0.003	0.022	0.004	0.023

A.2.5 Case F3.1: Flight Test 3, First Data Set . Case F3.1 was the first flight test of the October 2005 set. Based on the extensive problems during take-off in the April 2005 tests, the system was tested on the ground after engine start to verify proper operation, and then terminated prior to taxi. Once both aircraft were airborne, the system was re-engaged for testing. Also, for an unknown reason during the October flight tests, the floating filter's 1σ estimates of the relative position, relative velocity, and relative acceleration states were much higher than the April flight tests. When the April flight test data was post-processed using the code build from the October flight tests, the floating filter's 1σ estimates were similar to the October flight tests. This indicated that some code change impacted the filter's performance. There were no changes made to the author's code between the two flight test periods. However, an underlying matrix library and some peripheral code did change. Time did not permit the author to identify the exact cause, but the matrix library change is the most likely culprit.

Each sortie was scheduled for a two hour duration. The two hours included a small amount of taxi time as well as flight time. Each sortie was split into two one hour data runs. Since the first data run did not start until after the aircraft were airborne, most of the second hour data runs from the October tests were shorter than one hour. The weather during the 05 October test was clear, with light winds. The data run began at 0920L and the aircraft ranged from 13.0 to 1512.6 meters apart. Figure A.127 shows the minimum indicator North, East, Down relative positions.

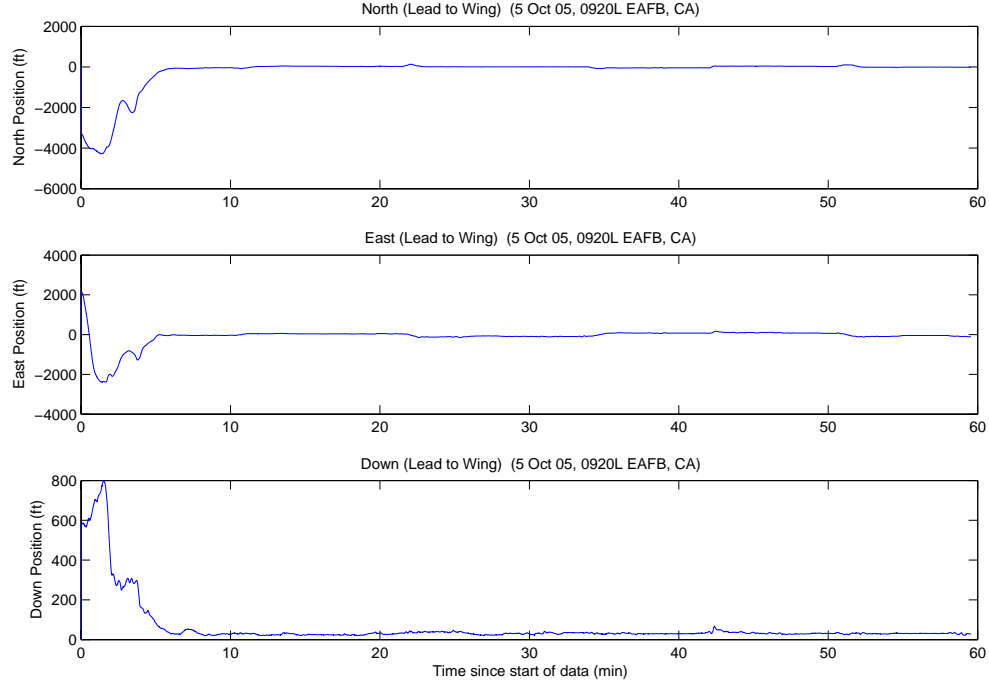


Figure A.127: Case F3.1: North, East, Down Relative Positions

The LAMBDA function returned a total of 28 unique ambiguity sets during the data run. The system was started once airborne when the aircraft were 1206 meters apart. Even with nine visible satellites, the true ambiguity set was the sixth set returned from the LAMBDA function; a full 50 seconds after the first ambiguity set was returned. The minimum indicator initially chose an incorrect ambiguity set at 4 minutes, and quickly changed to another incorrect set 20 seconds later. Finally, the minimum indicator selected the true ambiguity set at 4.8333 minutes, 30 seconds after the second incorrect set. After the minimum indicator selected the true set it did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 3378 of 3450 epochs (97.9 percent), with a maximum ratio of 28.7.

Figures A.128, A.129, and A.130 show the floating filter errors. The floating filter performed as expected based on initializing with nine satellites. The floating filter's 1σ estimate of the error is outside the plot scales in all three plots.

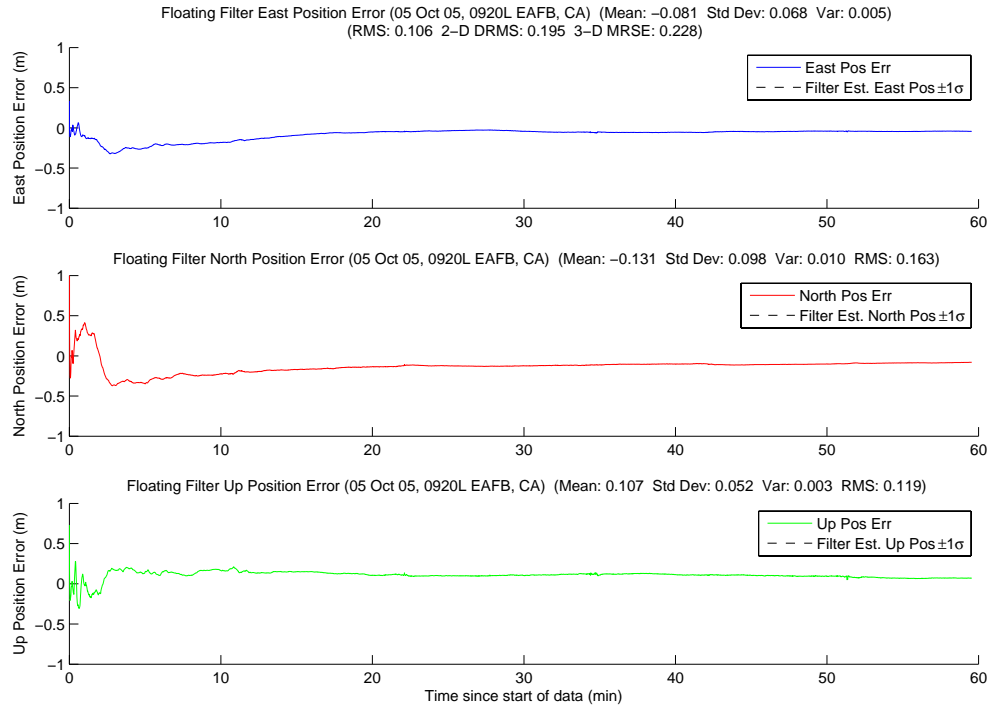


Figure A.128: Case F3.1: Floating Filter Position Errors

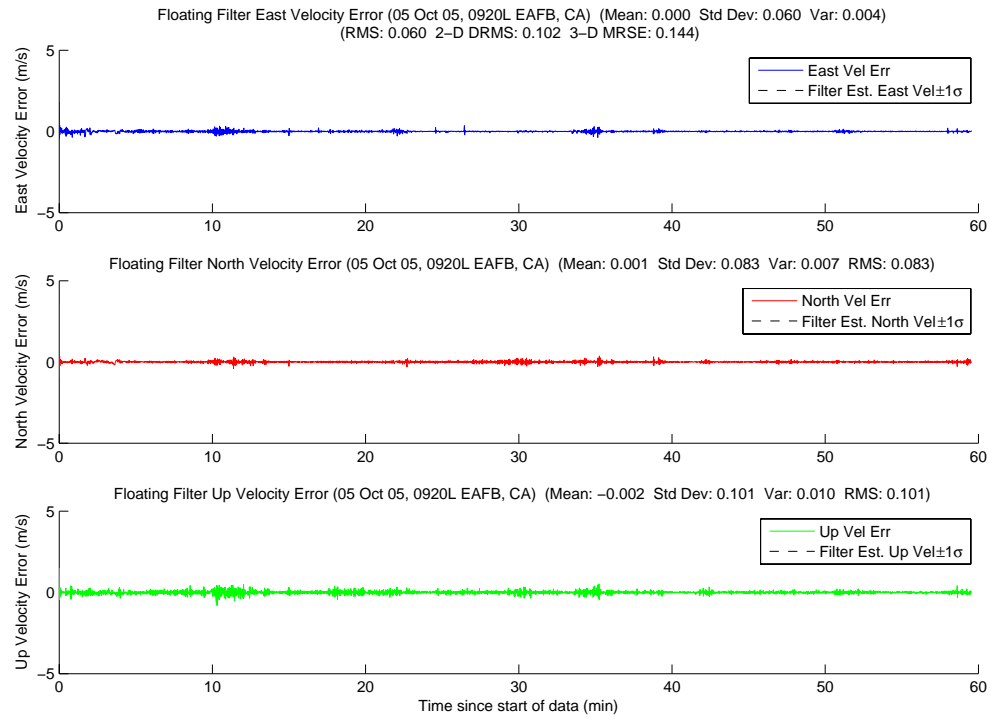


Figure A.129: Case F3.1: Floating Filter Velocity Errors

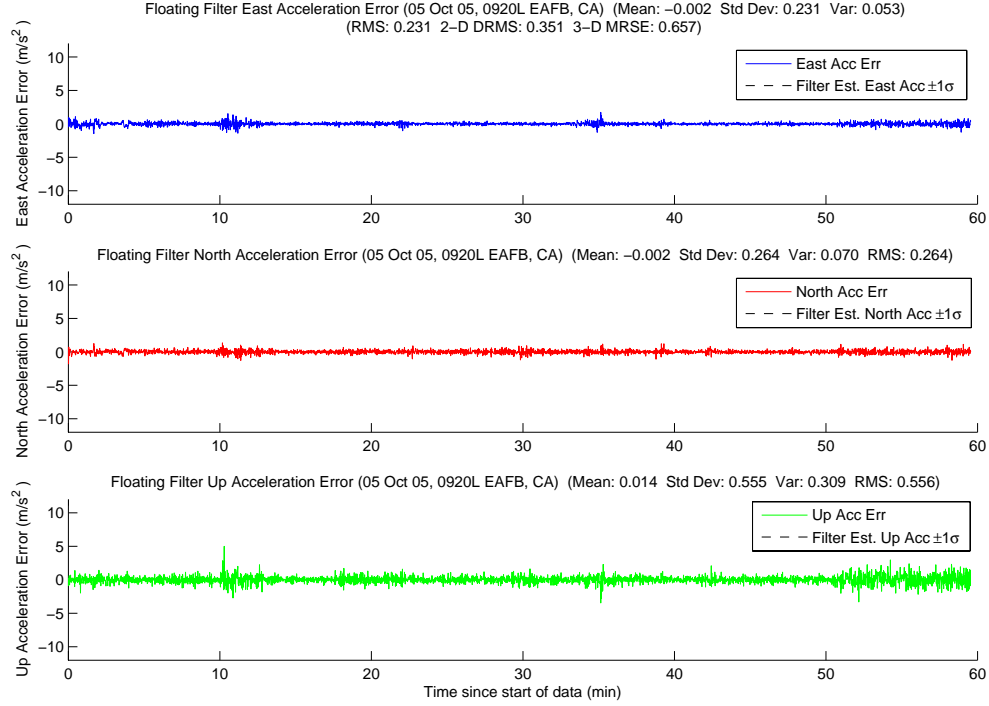


Figure A.130: Case F3.1: Floating Filter Acceleration Errors

Figure A.131 displays PRN 16 measurement information and the floating filter ambiguity estimate. The decrease in the floating filter 1σ ambiguity estimate at approximately 12 minutes was due to a change in the maximum elevated satellite. Otherwise, the plot is unremarkable. The rest of the satellite's plots follow in Figures A.132, A.133, A.134, A.135, A.136, A.137, A.138, A.139, and A.140.

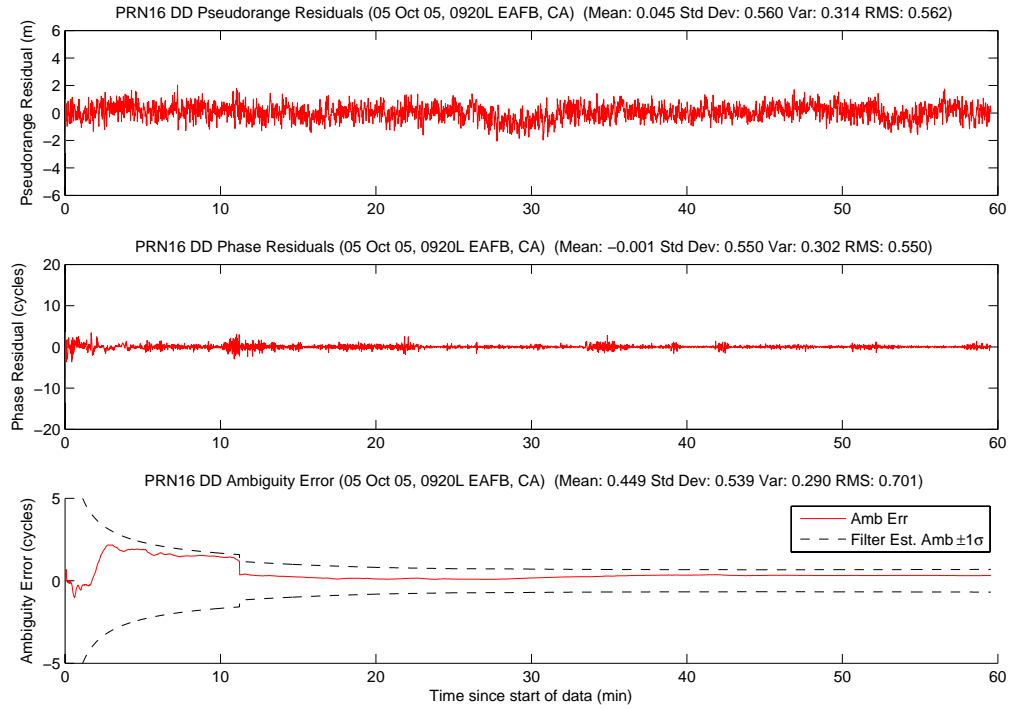


Figure A.131: Case F3.1: Satellite 16 Measurements

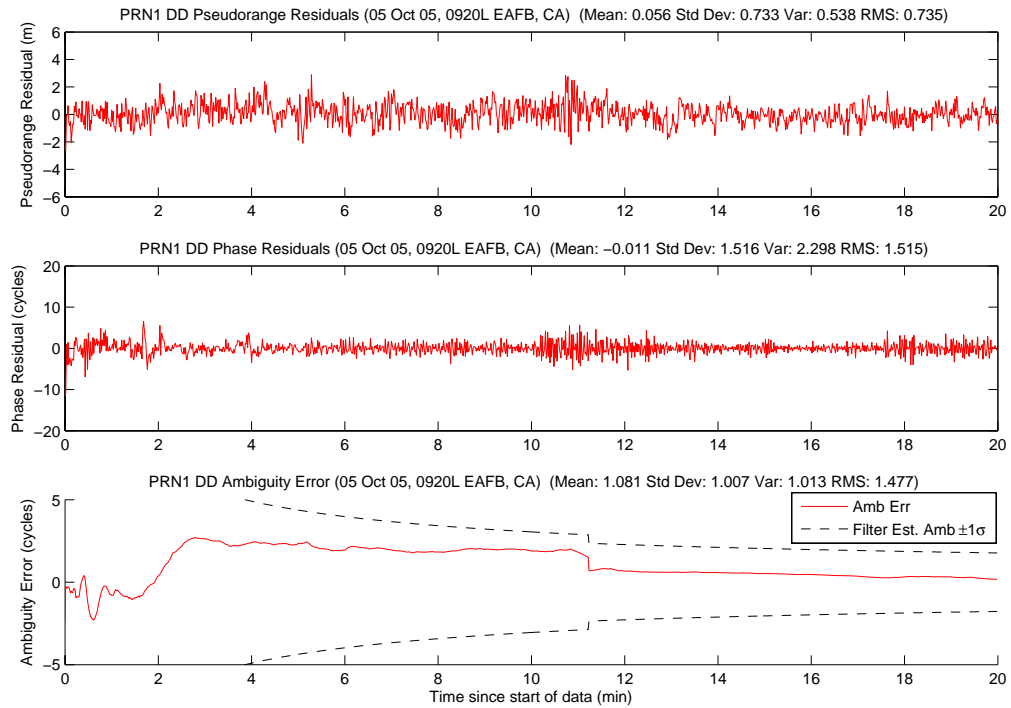


Figure A.132: Case F3.1: Satellite 1 Measurements

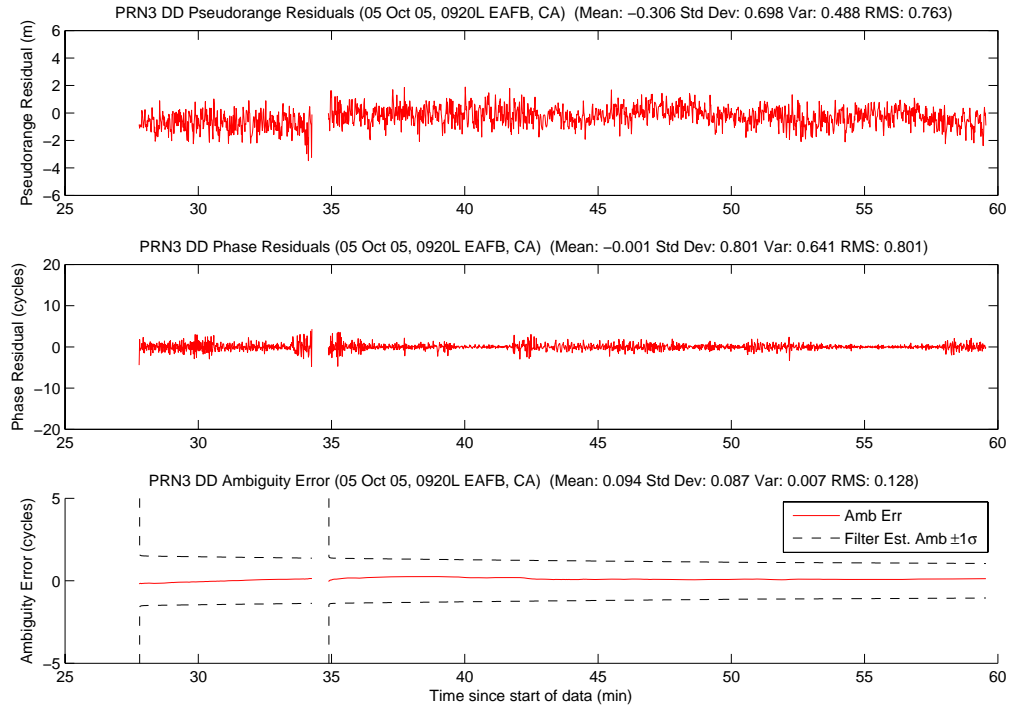


Figure A.133: Case F3.1: Satellite 3 Measurements

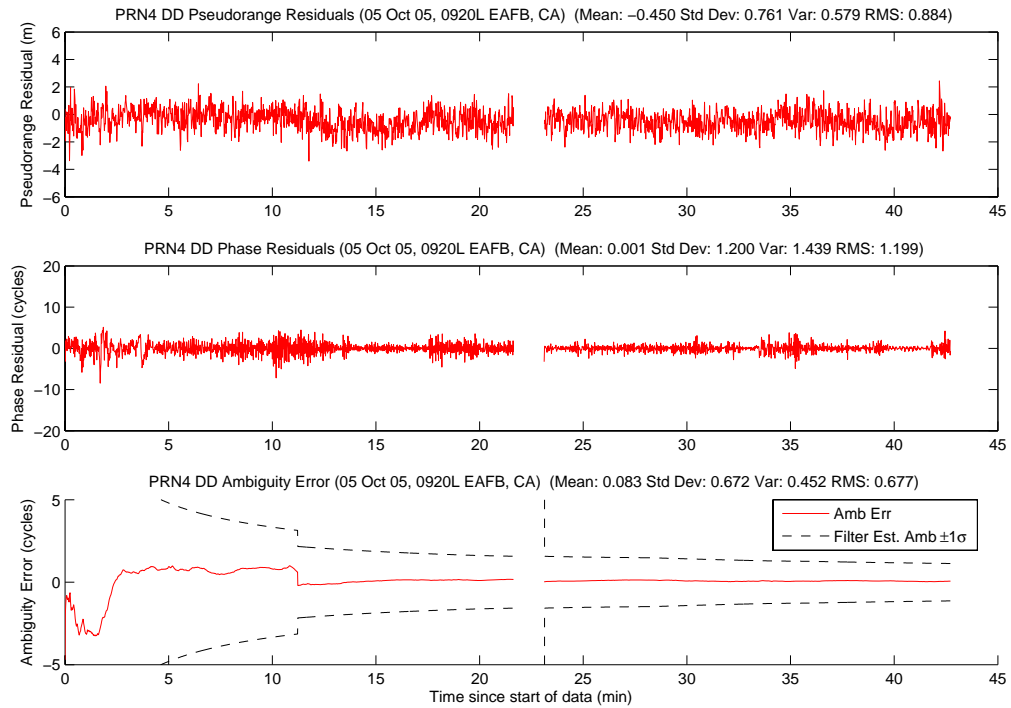


Figure A.134: Case F3.1: Satellite 4 Measurements

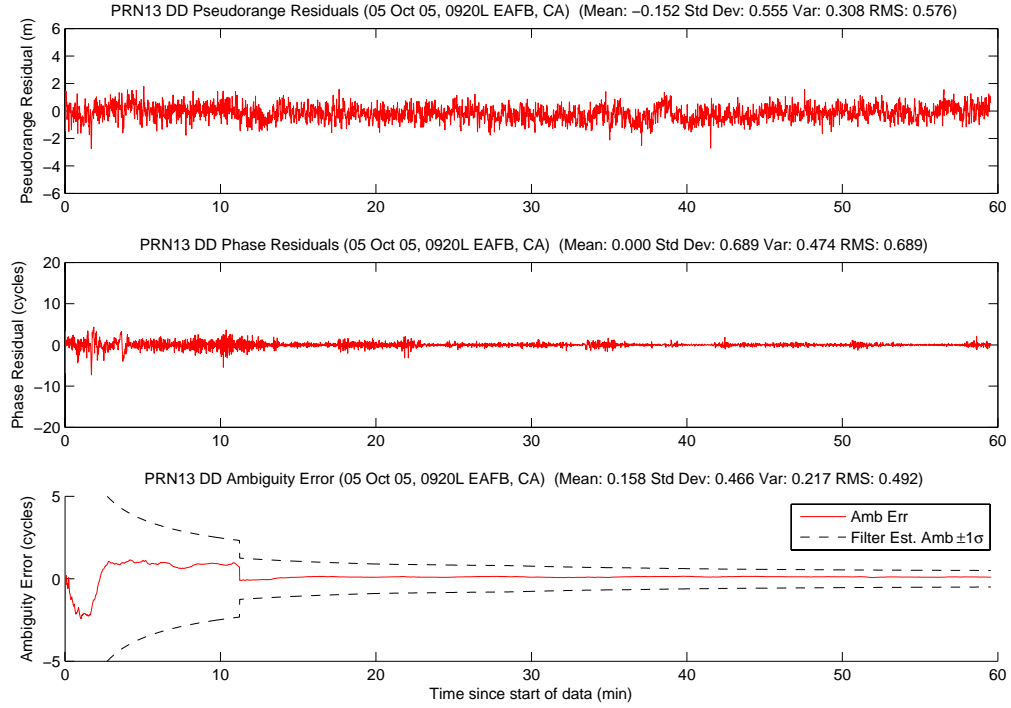


Figure A.135: Case F3.1: Satellite 13 Measurements

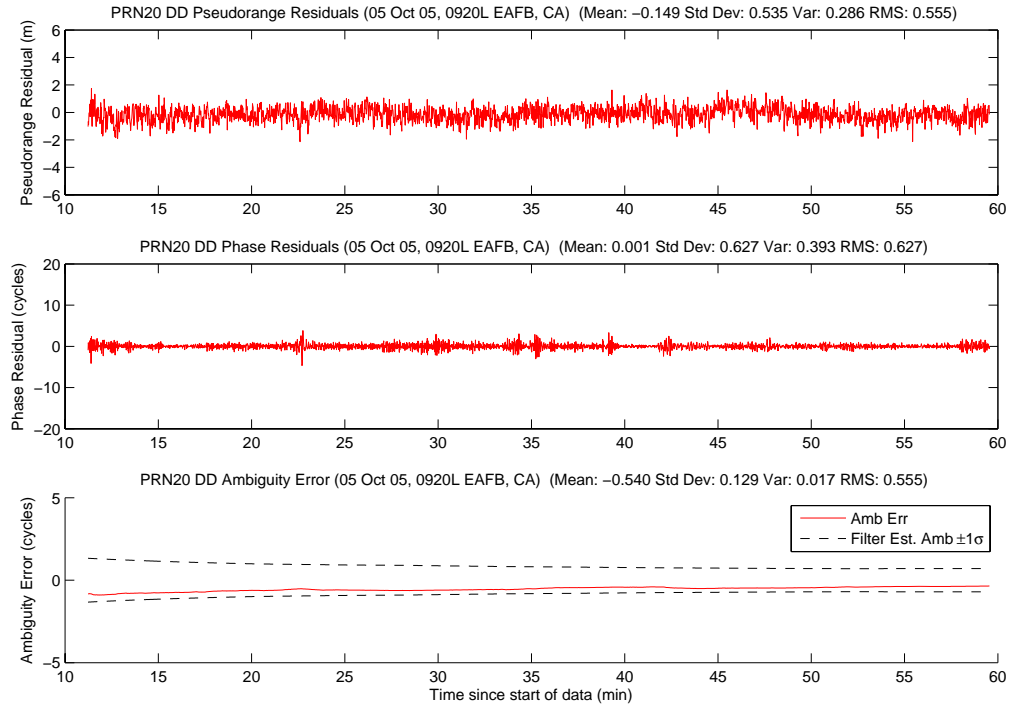


Figure A.136: Case F3.1: Satellite 20 Measurements

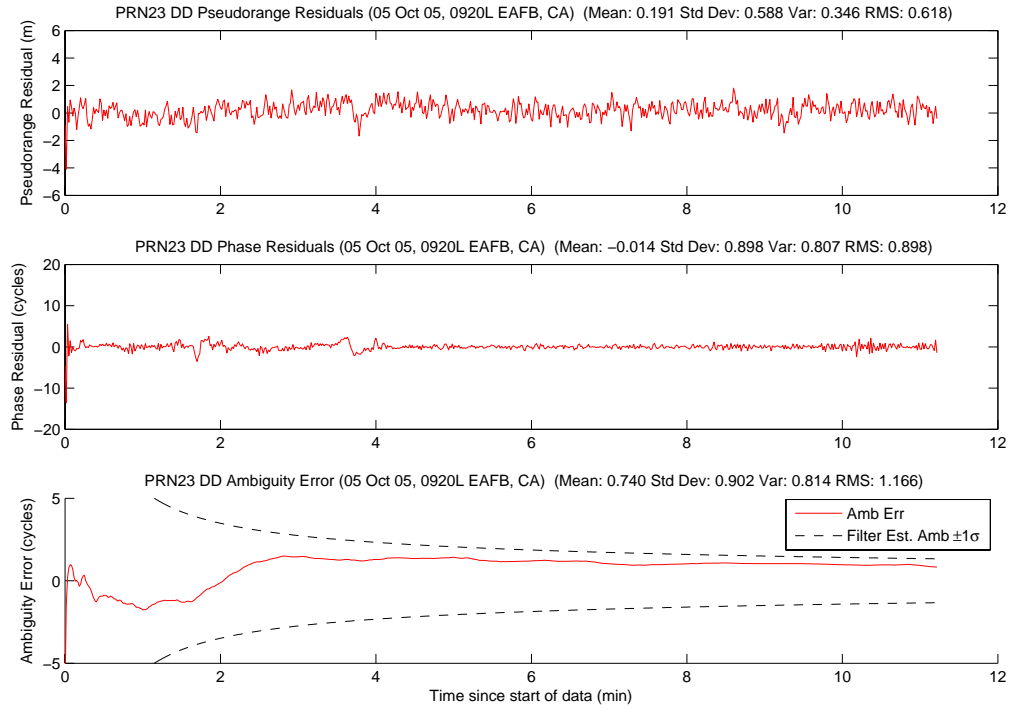


Figure A.137: Case F3.1: Satellite 23 Measurements

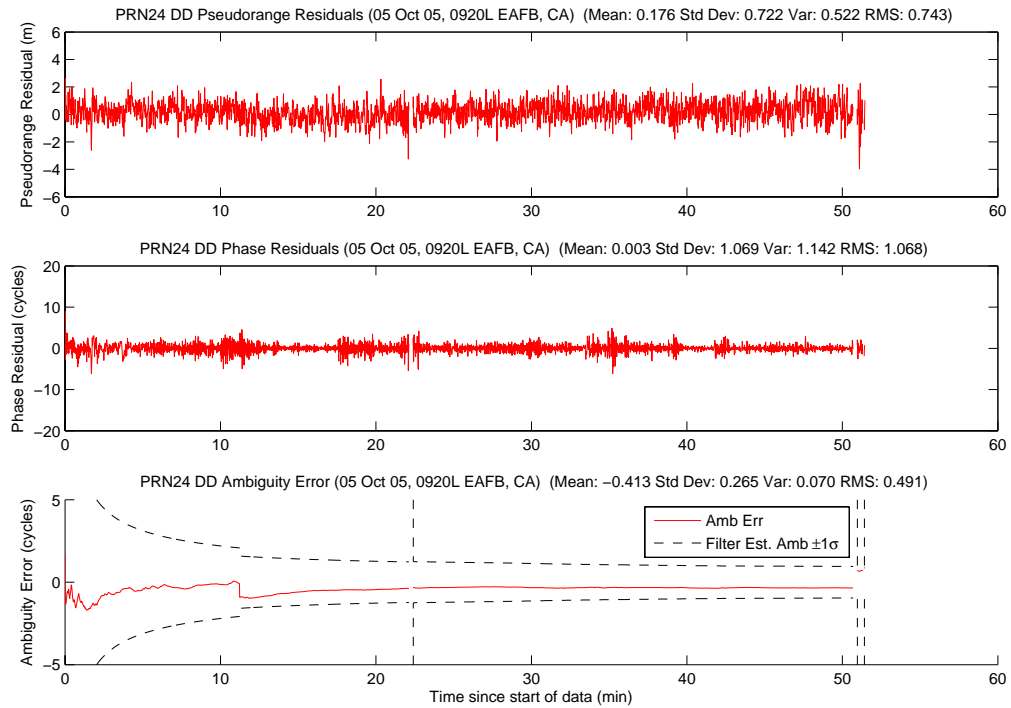


Figure A.138: Case F3.1: Satellite 24 Measurements

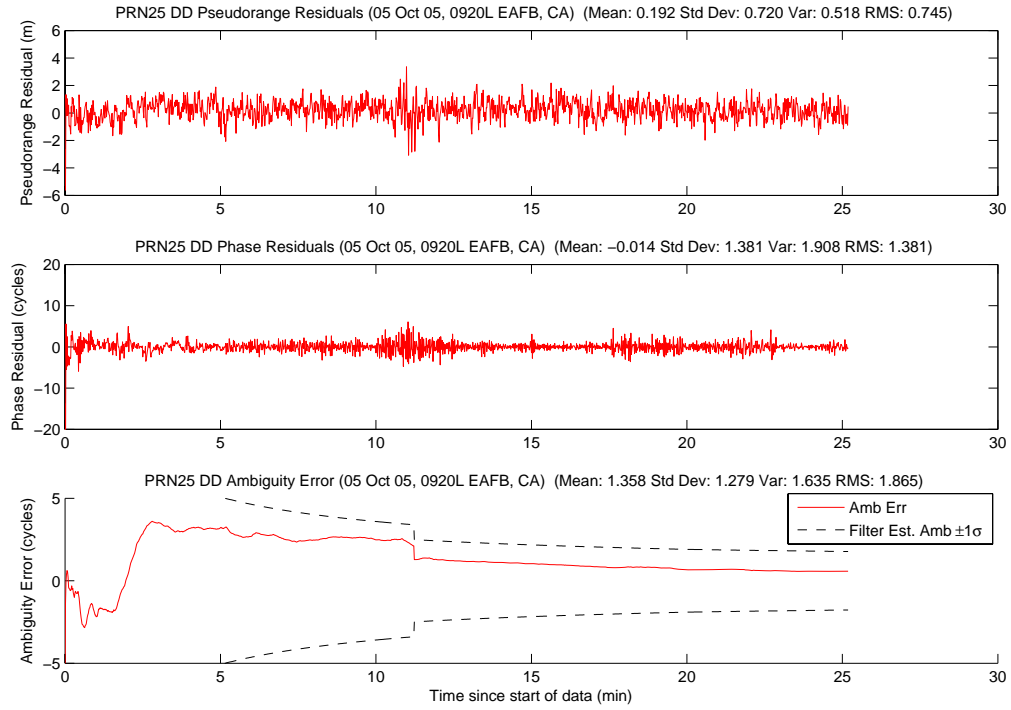


Figure A.139: Case F3.1: Satellite 25 Measurements

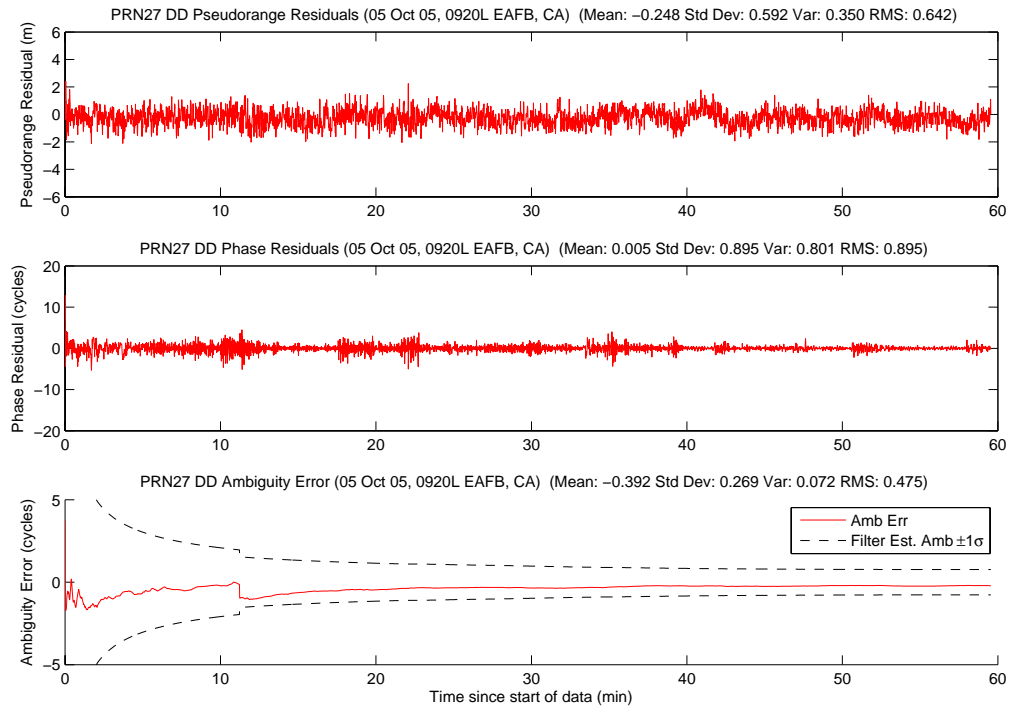


Figure A.140: Case F3.1: Satellite 27 Measurements

Figure A.141 shows the number of visible satellites. The satellite dropouts at 20, 42, and 50 minutes were due to setting satellites (PRN 1, PRN 4, and PRN 24 respectively). The other dropouts at 22, 23, 25, and 32 minutes were due to satellite masking. Finally, the small dropout at 49 minutes was a cycle slip on PRN 24 just prior to the satellite setting below 10 degrees.

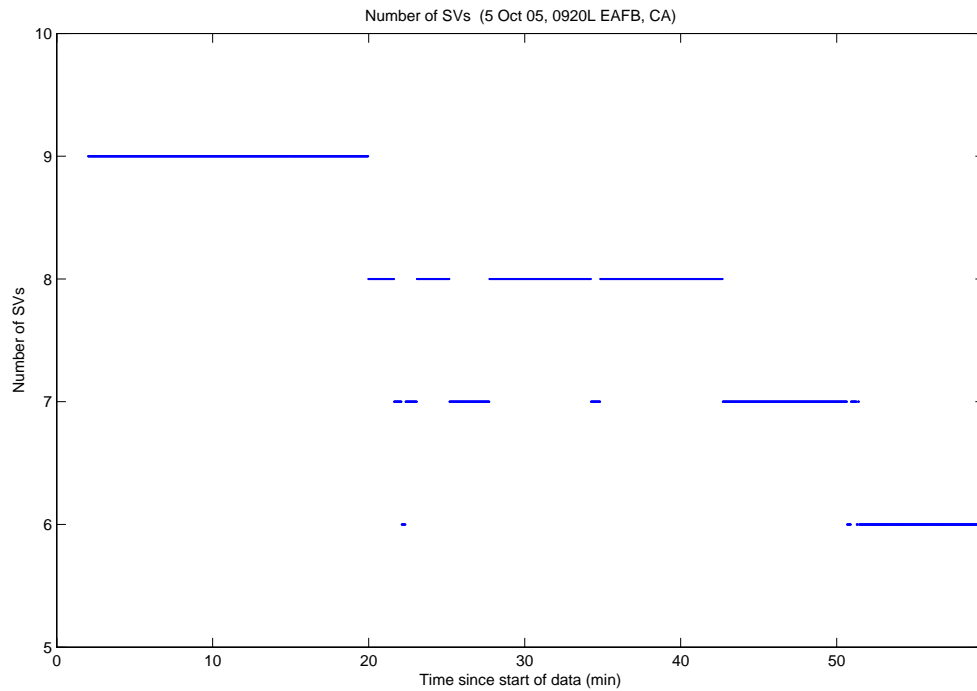


Figure A.141: Case F3.1: Visible Satellites

Figure A.142 shows all of the candidate ambiguity set's SOS residuals. It is interesting to note the slow increase and then decrease of every candidate set's SOS residuals starting at 36 minutes until 44 minutes. Looking at the visible satellite plot reveals the possible culprit. PRN 4 went below the satellite cutoff angle at 42 minutes. Lower elevated satellites are effected more by the atmosphere and also more prone to effects such as multipath and masking. It was highly likely that PRN 4 measurements were corrupting the position solution. Once PRN 4 set, the residuals quickly returned to their original state. Also, erroneous set's residuals were below the true set at 11, 26, 49, 51, and 55 minutes. Figure A.143 displays the MMAE conditional probability. The true ambiguity set maintained over 49 percent of the total probability starting at 5 minutes into the data run. This was approximately the same time the minimum indicator selected the true ambiguity set. The probability drops at 11, 26, 49, 51, and 55 minutes correlate directly to the SOS residual plot.

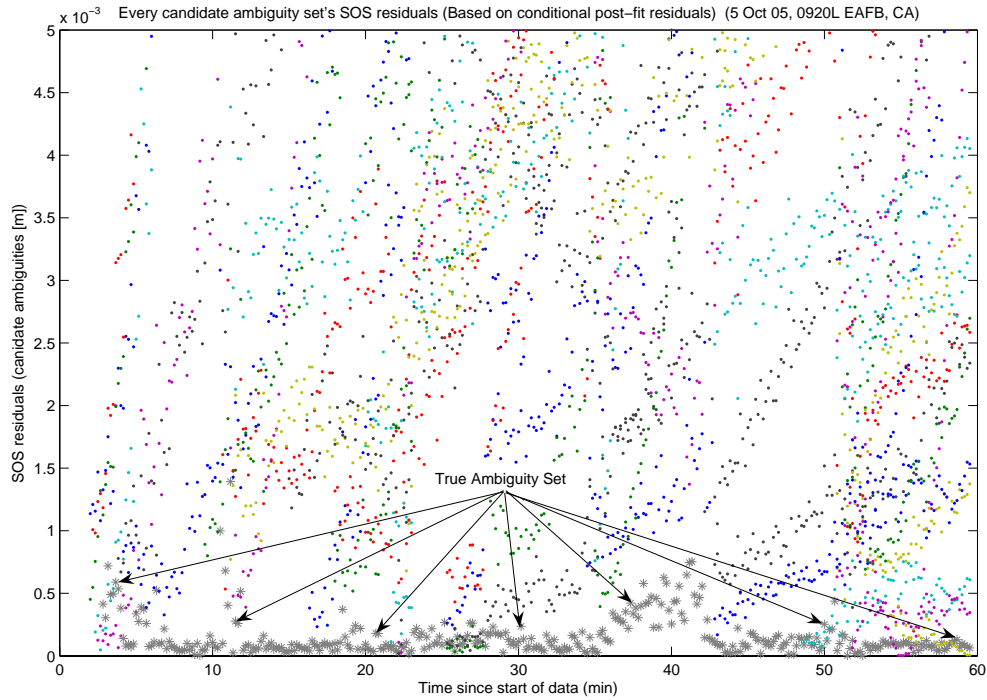


Figure A.142: Case F3.1: SOS Residuals

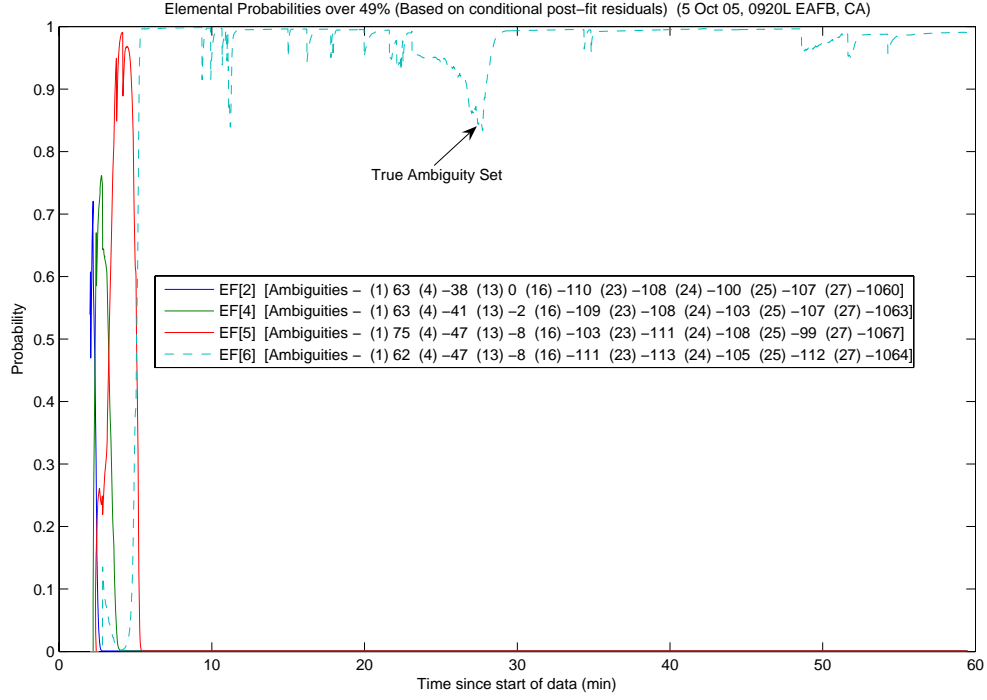


Figure A.143: Case F3.1: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.144 is EF[2], Figure A.145 is EF[4], Figure A.146 is EF[5], and Figure A.147 is EF[6].

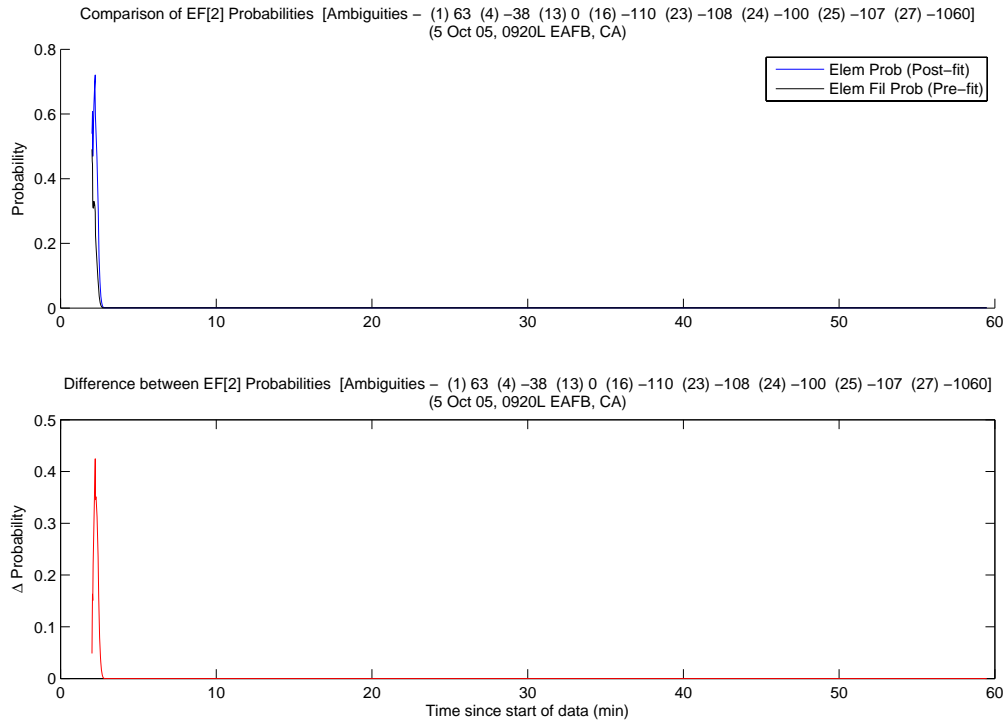


Figure A.144: Case F3.1: EF[2] Probability Comparison

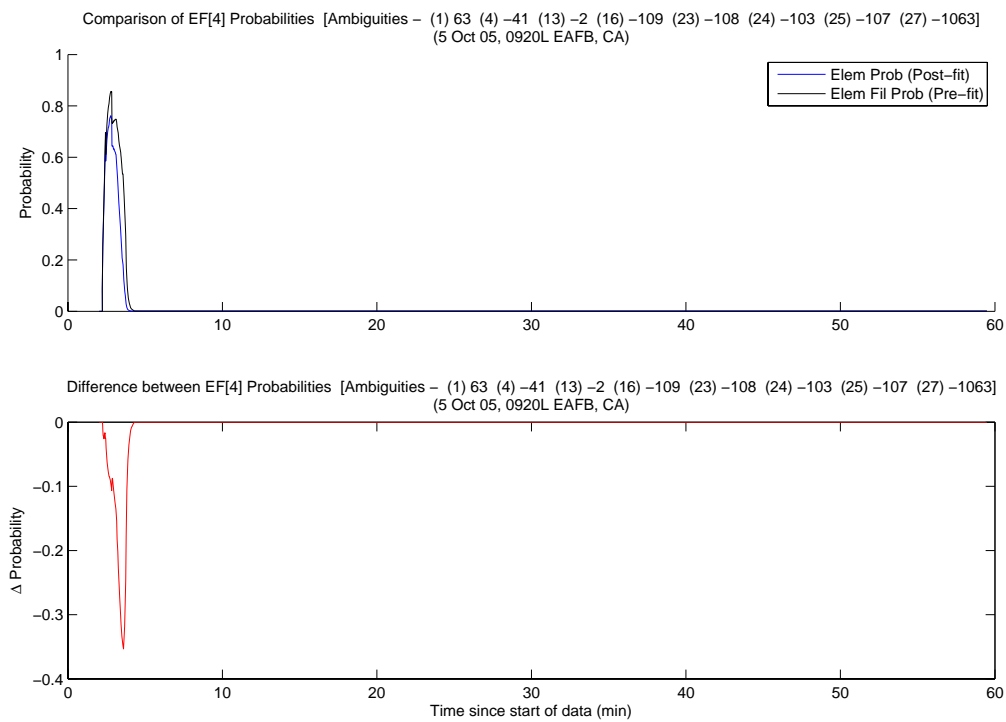


Figure A.145: Case F3.1: EF[4] Probability Comparison

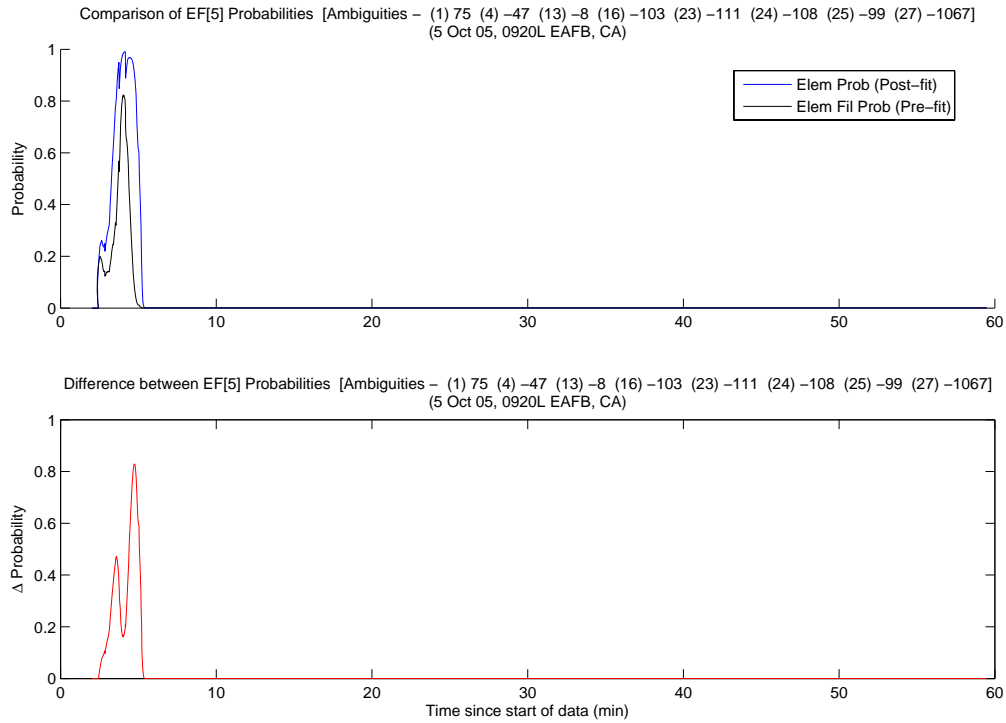


Figure A.146: Case F3.1: EF[5] Probability Comparison

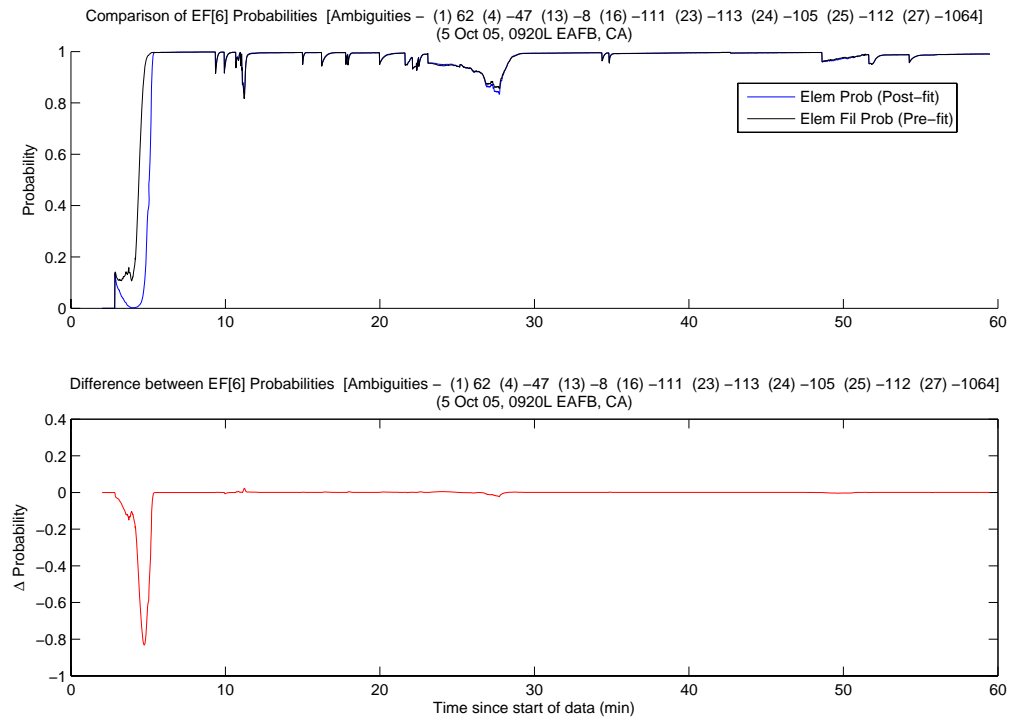


Figure A.147: Case F3.1: EF[6] Probability Comparison

The MMAE position error is shown in Figure A.148. Initially, the MMAE solution was incorrect due to three erroneous ambiguity sets absorbing the bulk of the probability as seen in Figure A.143. The smaller effects from the other probability dips are not visible because of the scale.

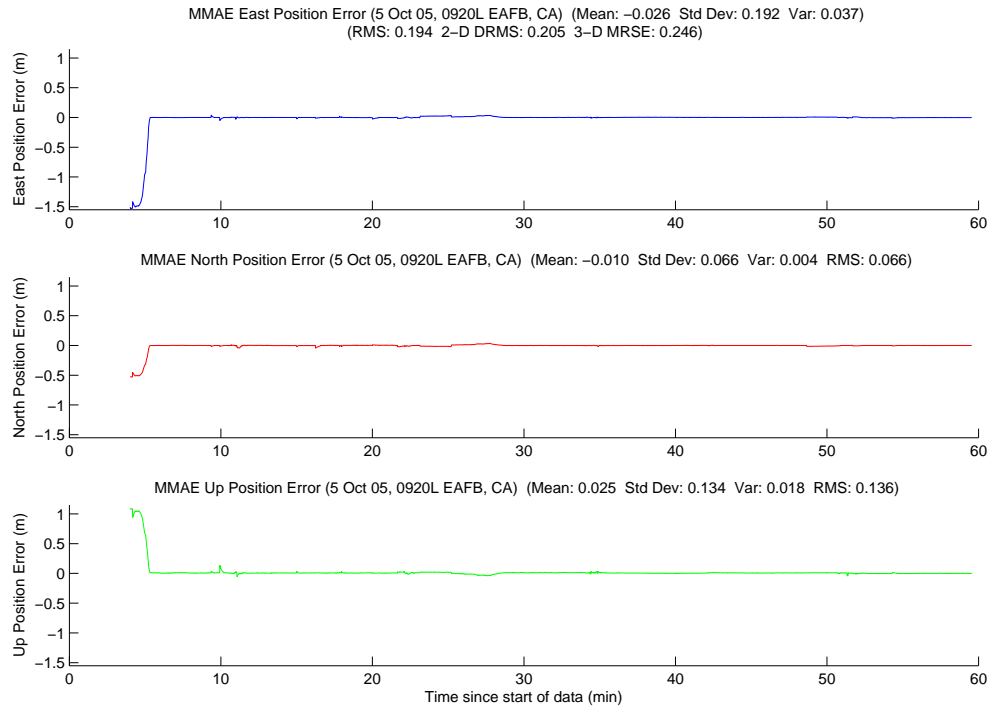


Figure A.148: Case F3.1: MMAE Position Error

Finally, Figure A.149 displays the position error plot for the minimum indicator variable. Unfortunately, since the minimum indicator incorrectly chose two other sets before finally selecting the true ambiguity set, the MRSE was outside the centimeter-level requirement. However, the aircraft were 1200 meters apart when the error from the minimum indicator was at it highest, 1.5 meters. Once the aircraft joined up, the position solution from the minimum indicator was well within the centimeter-level requirement for the remainder of the flight. Lastly, Table A.9 gives a summary of each method's position error.

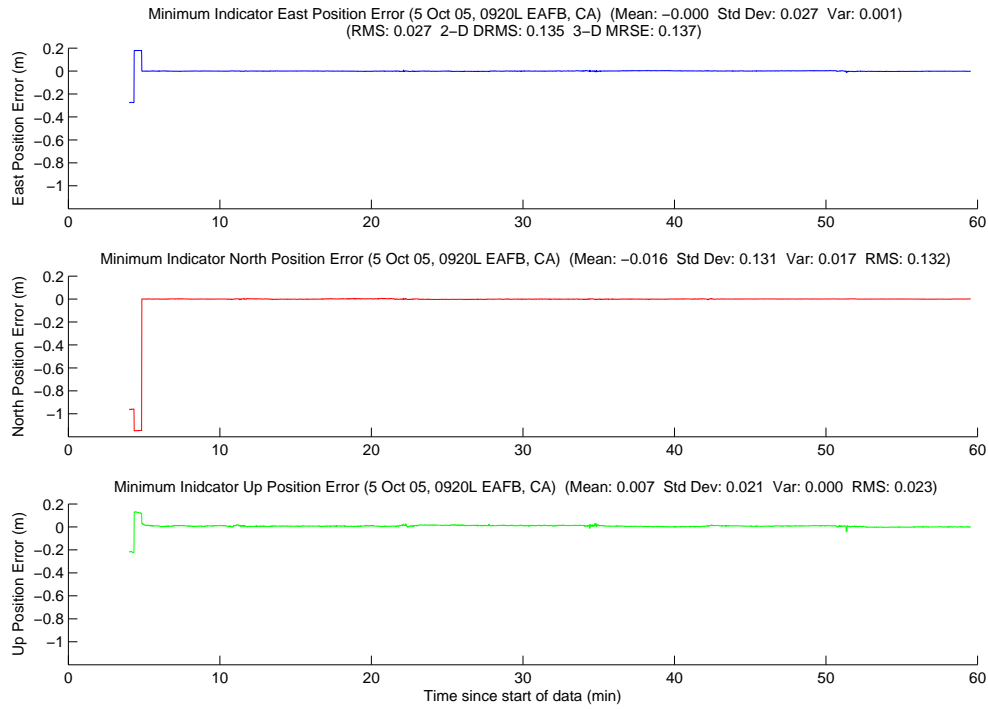


Figure A.149: Case F3.1: Minimum Indicator Position Error

Table A.9: Case F3.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.081	0.068	-0.131	0.098	0.107	0.052	0.195	0.228
MMAE	-0.026	0.192	-0.010	0.066	0.025	0.134	0.205	0.246
Min. Ind.	-0.000	0.027	-0.016	0.131	0.007	0.021	0.135	0.137

A.2.6 Case F3.2: Flight Test 3, Second Data Set. Case F3.2 started within a minute of the end of Case F3.1. The data run began at 1020L and the aircraft ranged from 13.0 to 1512.6 meters apart. Figure A.150 shows the minimum indicator North, East, Down relative positions. The data run ended at 37 minutes since the aircraft had separated for return to base (RTB).

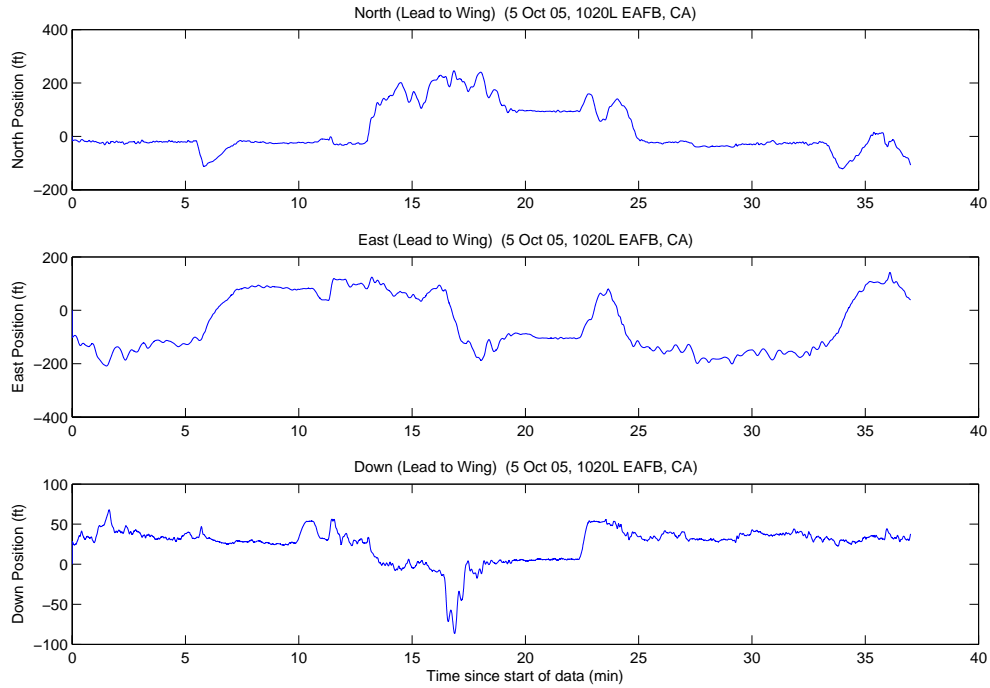


Figure A.150: Case F3.2: North, East, Down Relative Positions

Since the data set was terminated early, some information about the run was unavailable. The LAMBDA function returned over 50 unique ambiguity sets during the data run. The system started when the aircraft were 32.7 meters apart, with six visible satellites. The true ambiguity set was the second set, of the first pair, returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at 4 minutes. The minimum indicator remained “locked” on the true set until 13 minutes, when the filter was reset due to missing messages and checksum failures. For an unknown reason the datalink was having problems communicating, even though the aircraft were only 27 meters apart. Next, at 15.5 minutes, 2.5 minutes after the reset, the minimum indicator chose an erroneous ambiguity set. Then 40 seconds later, the minimum indicator chose the true ambiguity set, which was the fifth set returned from the LAMBDA function after the reset. Interestingly, the ambiguities for the second true set were the same as the first true set. Unfortunately, 14 seconds later, the datalink began to have problems again with the aircraft 70 meters apart. The relative position module erroneously detected four cycle slips, one of which was on the base satellite, and the filter was reset again. The minimum indicator found the true ambiguity set, for the third time, at 18.5 minutes. This was two minutes after the second filter reset. The third true ambiguity set was the first set returned from the LAMBDA function after the second reset. The ambiguities were once again the same as the previous two true ambiguity sets. The minimum indicator remained “locked” to the true set for the remainder of the data run. The true set was returned from LAMBDA as number one, 2021 of 2220 epochs (91.0 percent), even with the filter resets.

Figures A.151, A.152, and A.153 show the floating filter errors. The two filter resets at 13 and 16 minutes are visible in the relative position plot as jumps in the floating filter position errors. The floating filter’s 1σ estimate of the error is outside all three plot scales.

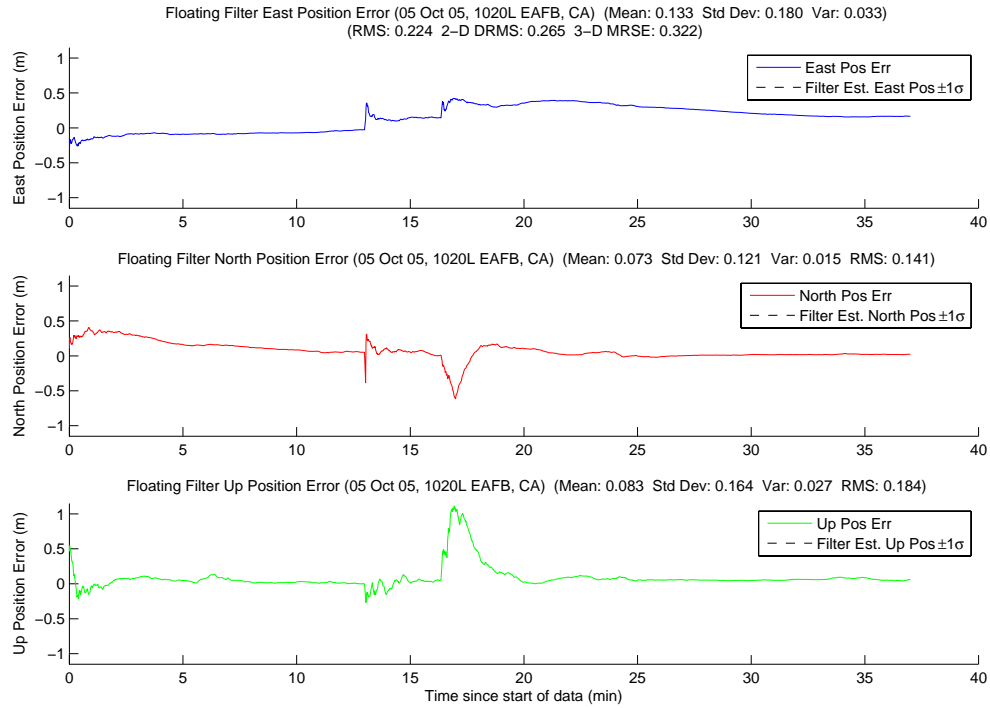


Figure A.151: Case F3.2: Floating Filter Position Errors

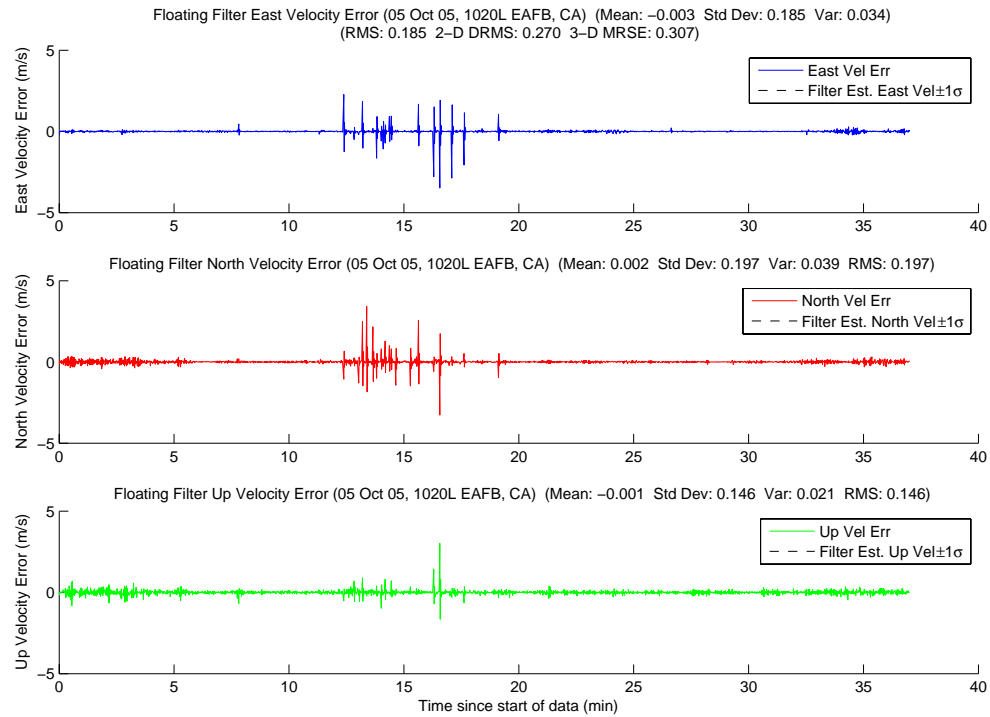


Figure A.152: Case F3.2: Floating Filter Velocity Errors

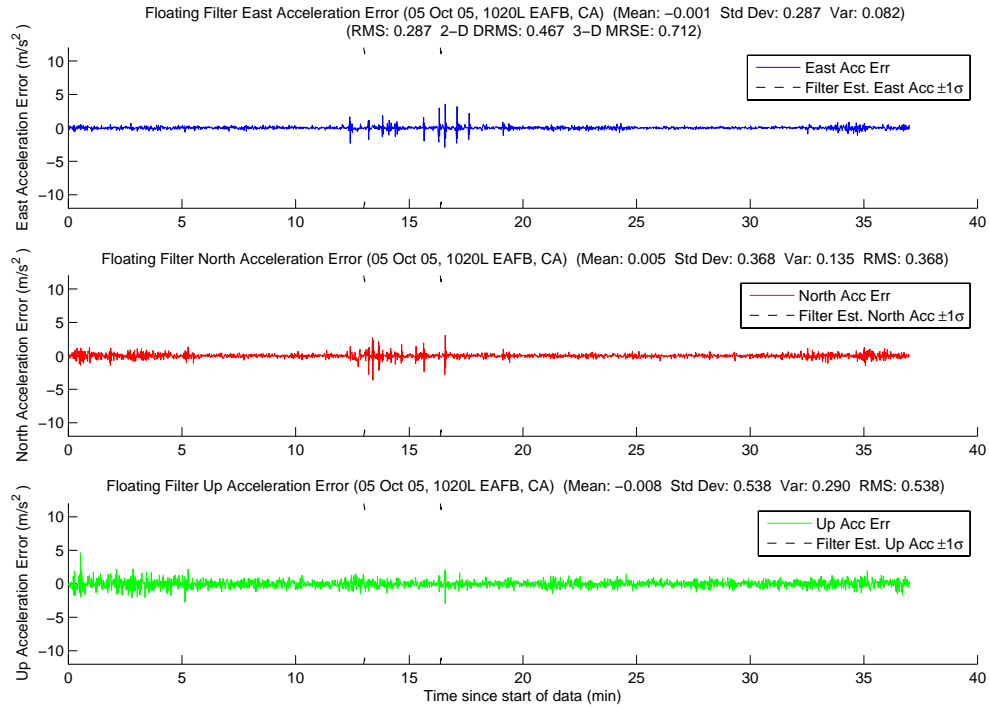


Figure A.153: Case F3.2: Floating Filter Acceleration Errors

Figure A.154 displays PRN 16 measurement information and the floating filter ambiguity estimate. The filter resets are once more visible in the plot. Otherwise, the plot is unremarkable. The rest of the satellite's plots follow in Figures A.155, A.156, A.157, A.158, A.159, A.160, and A.161.

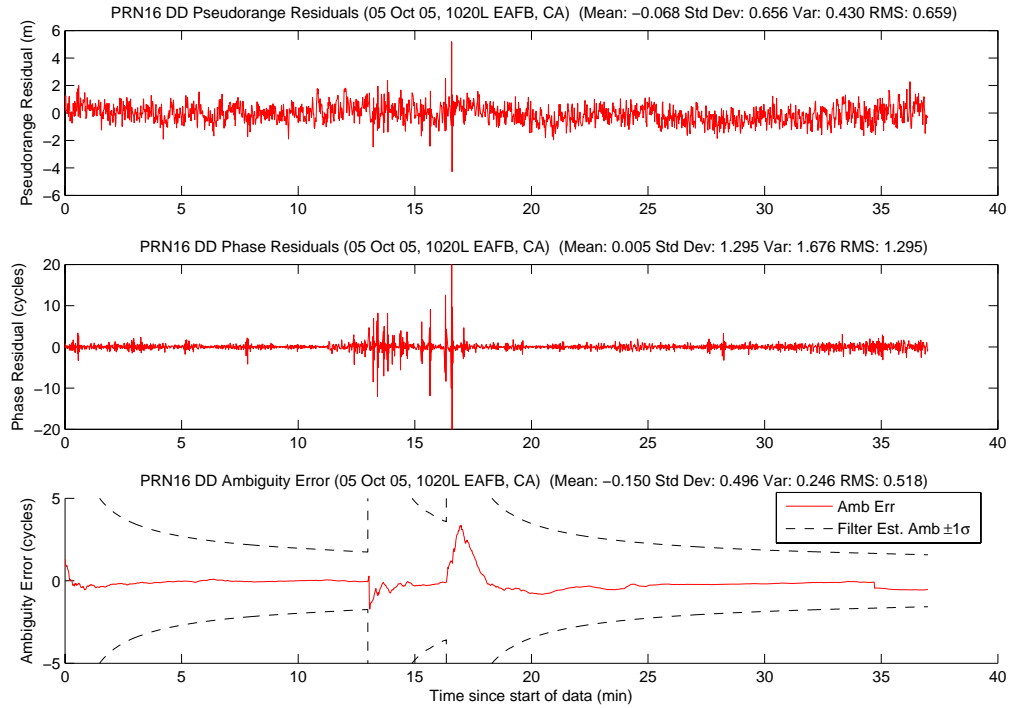


Figure A.154: Case F3.2: Satellite 16 Measurements

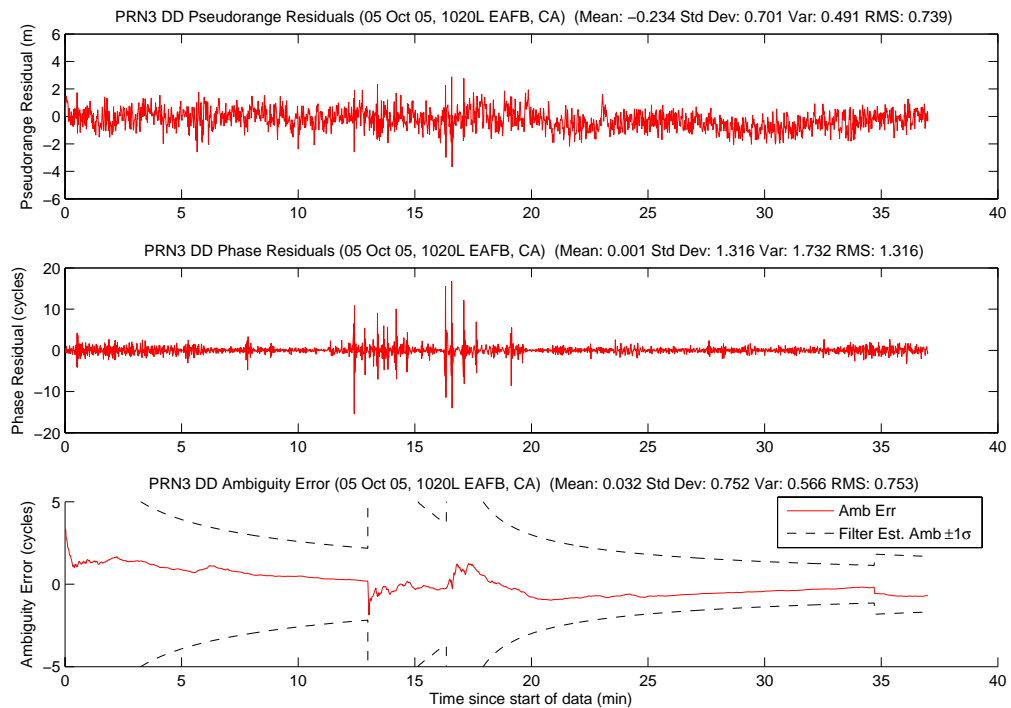


Figure A.155: Case F3.2: Satellite 3 Measurements

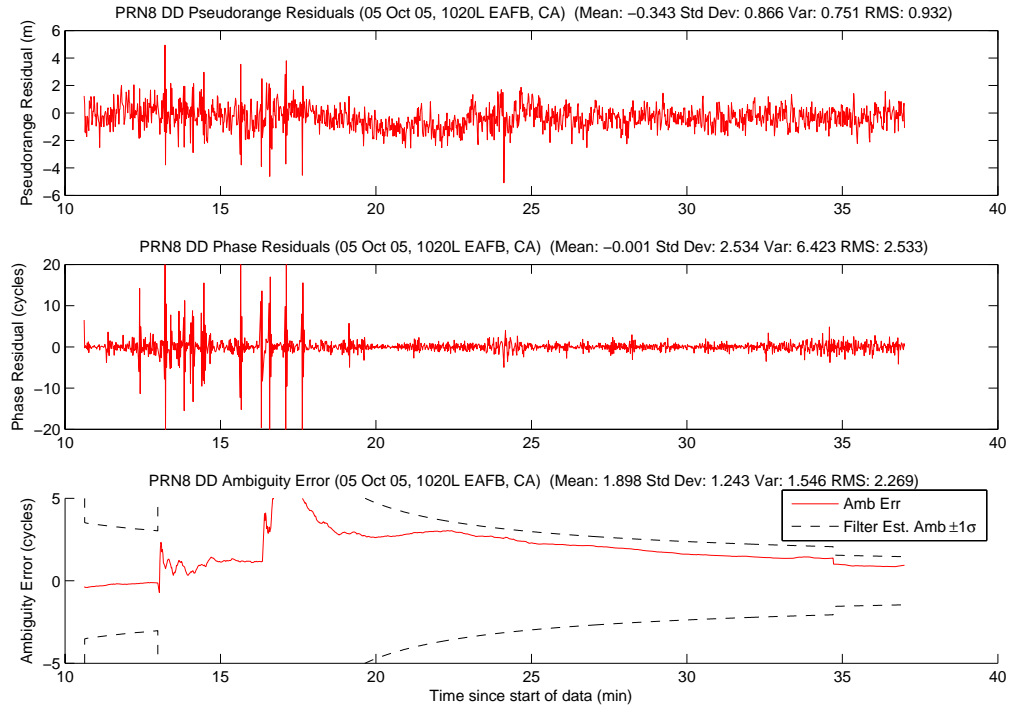


Figure A.156: Case F3.2: Satellite 8 Measurements

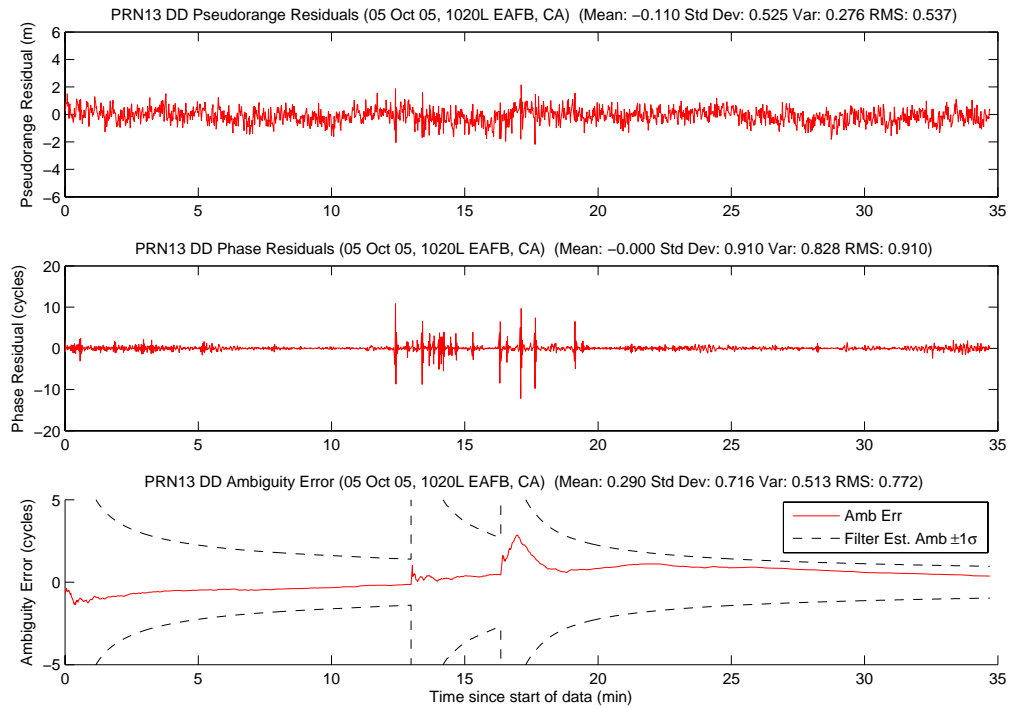


Figure A.157: Case F3.2: Satellite 13 Measurements

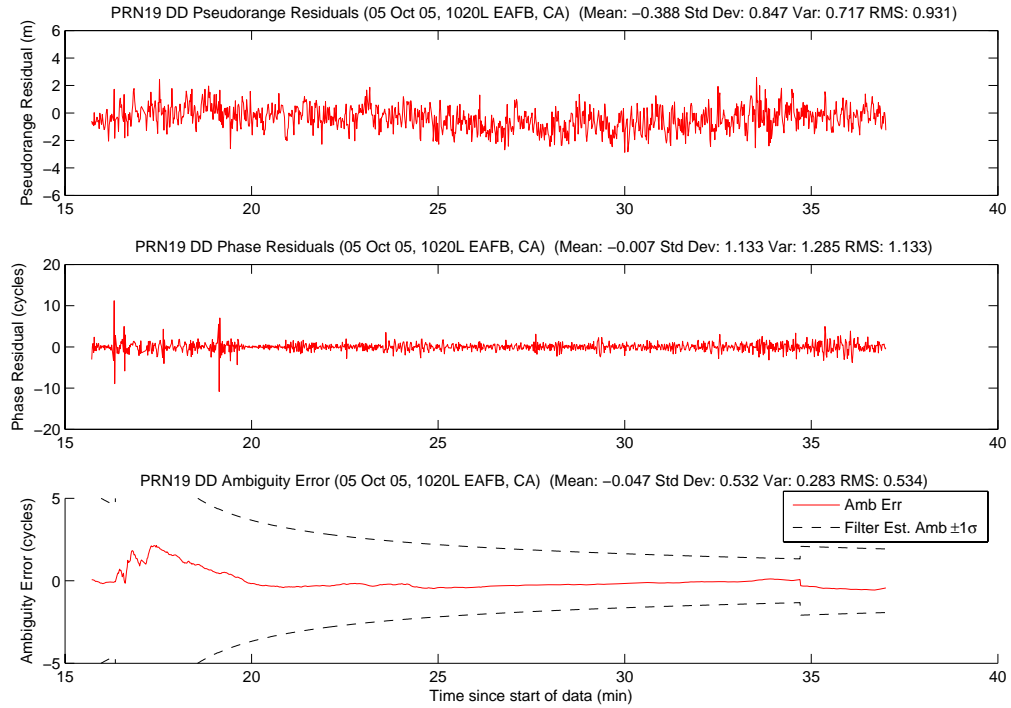


Figure A.158: Case F3.2: Satellite 19 Measurements

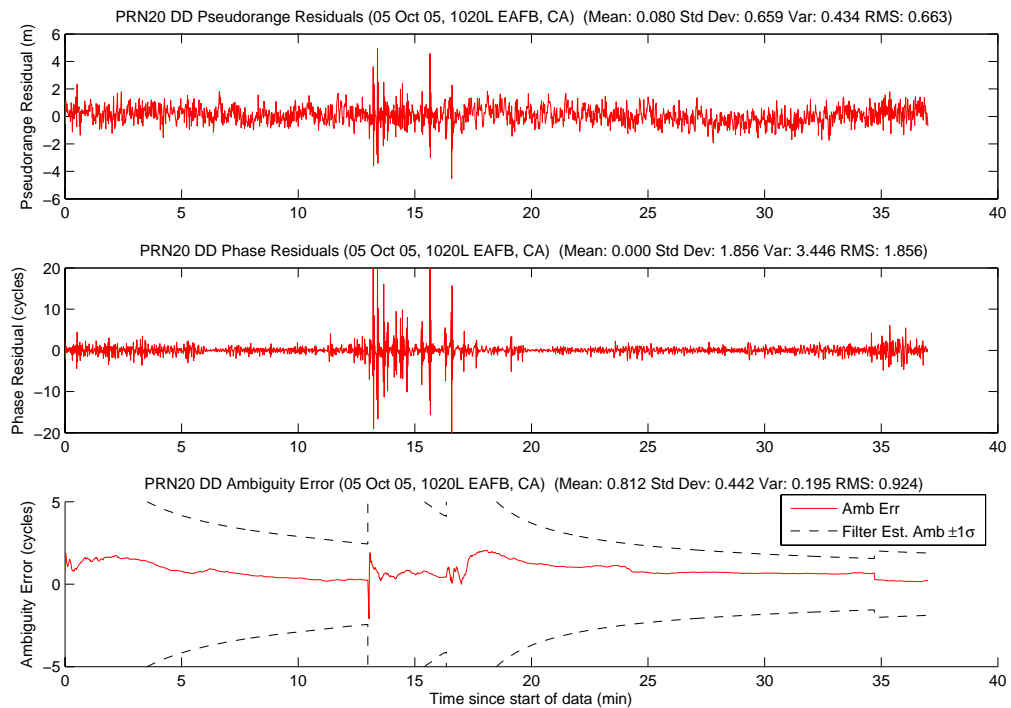


Figure A.159: Case F3.2: Satellite 20 Measurements

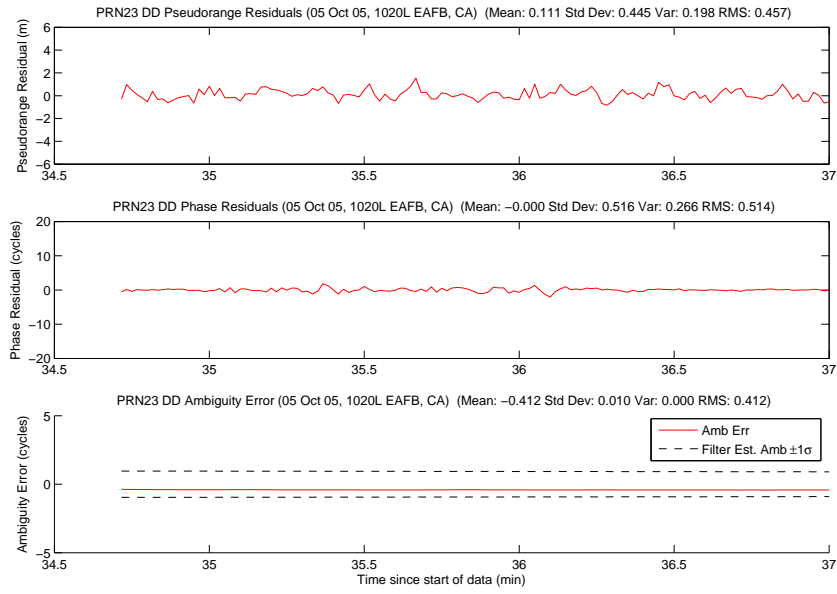


Figure A.160: Case F3.2: Satellite 23 Measurements

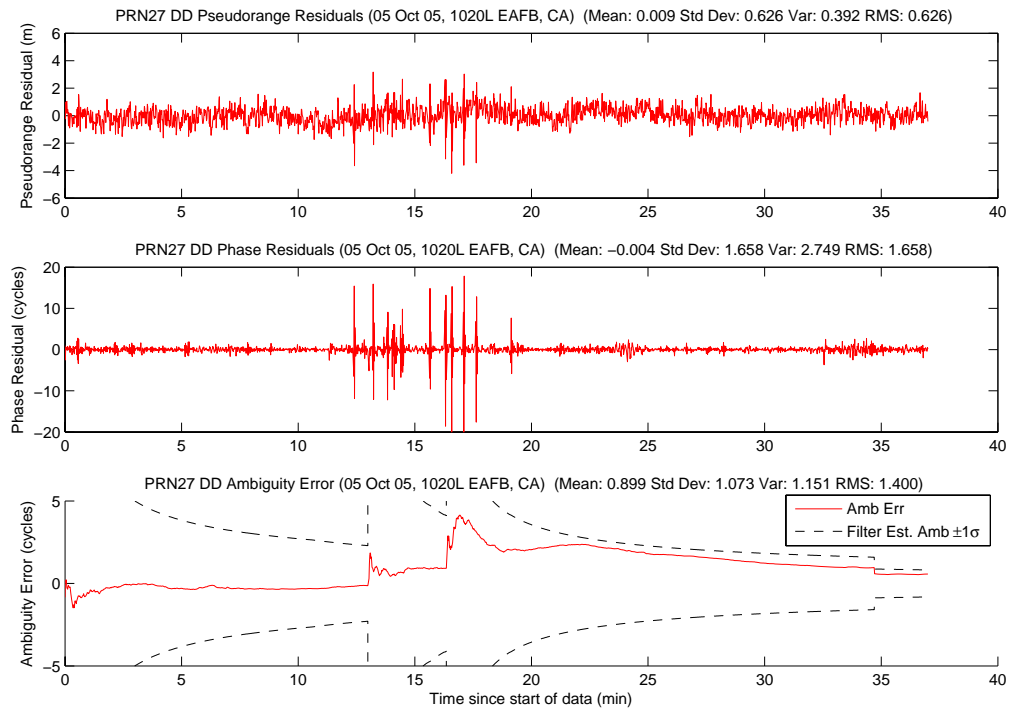


Figure A.161: Case F3.2: Satellite 27 Measurements

Figure A.162 shows the number of visible satellites. The satellite dropout at 16 minutes was from an erroneously detected cycle slip, as mention earlier.

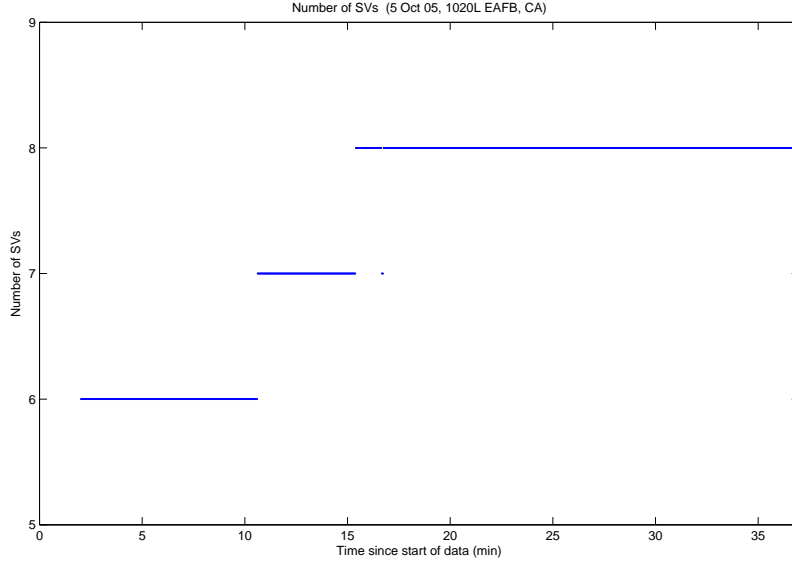


Figure A.162: Case F3.2: Visible Satellites

Figure A.163 shows all of the candidate ambiguity set's SOS residuals, and Figure A.164 displays the MMAE conditional probability. The three true ambiguity sets are clearly visible in probability plot. At 11 minutes, before the first filter reset, the true ambiguity set's probability sharply declined as EF[7] absorbed over 50 percent of the probability. Looking at the SOS residual plot, Figure A.163, reveals the reason. An erroneous set's residuals (black dots on the plot) were well below the true ambiguity set starting at 11 minutes until the first filter reset. Also, the other probability dip at 22 minutes was due to an erroneous set's residuals approaching the true set's residuals (gold dots on the plot).

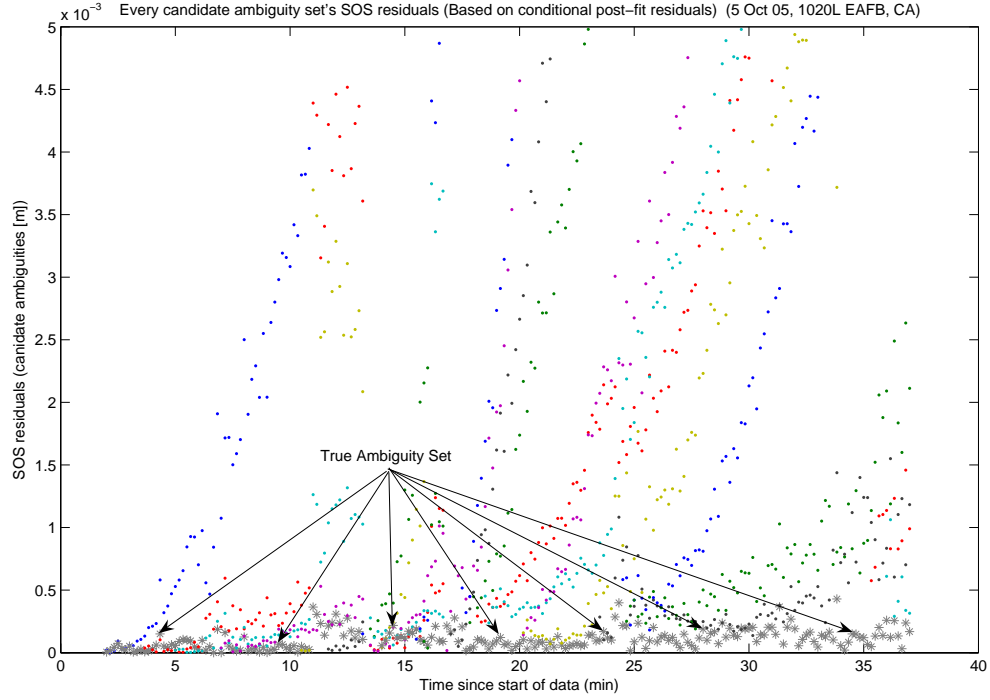


Figure A.163: Case F3.2: SOS Residuals

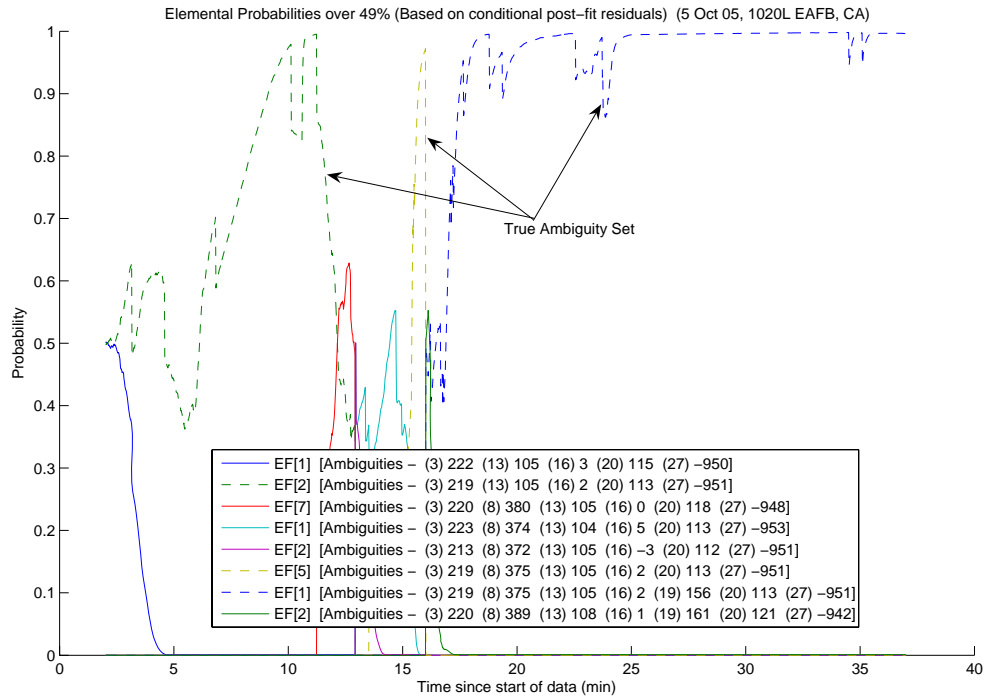


Figure A.164: Case F3.2: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.165 is EF[1], Figure A.166 is EF[1A], Figure A.167 is EF[1B], Figure A.168 is EF[2], Figure A.169 is EF[2A], Figure A.170 is EF[2B], and Figure A.171 is EF[5].

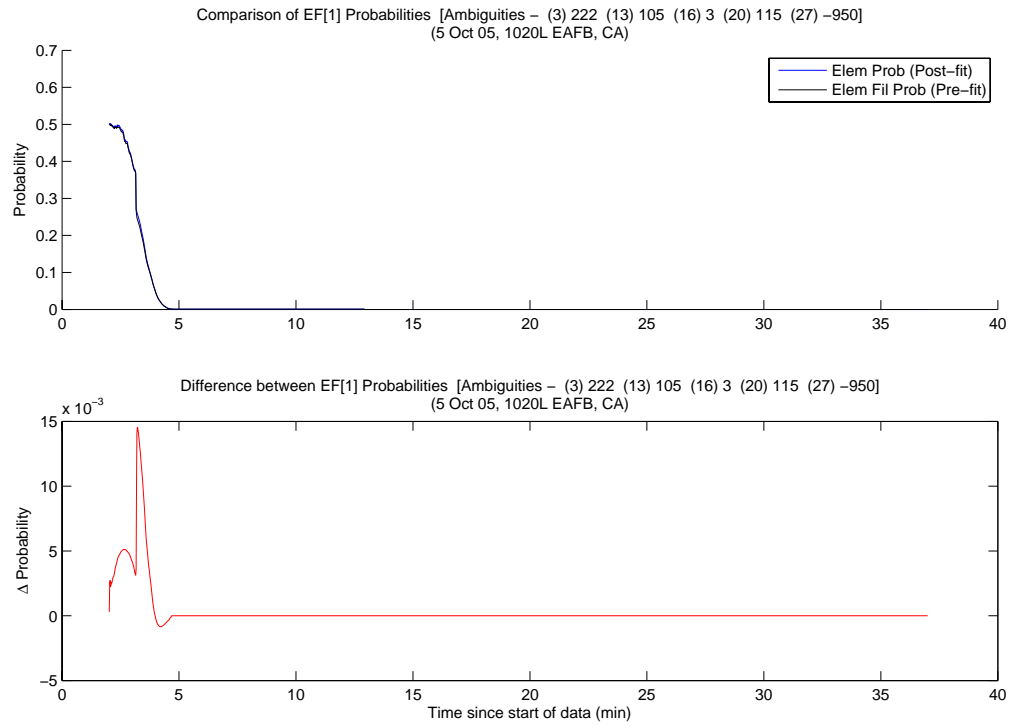


Figure A.165: Case F3.2: EF[1] Probability Comparison

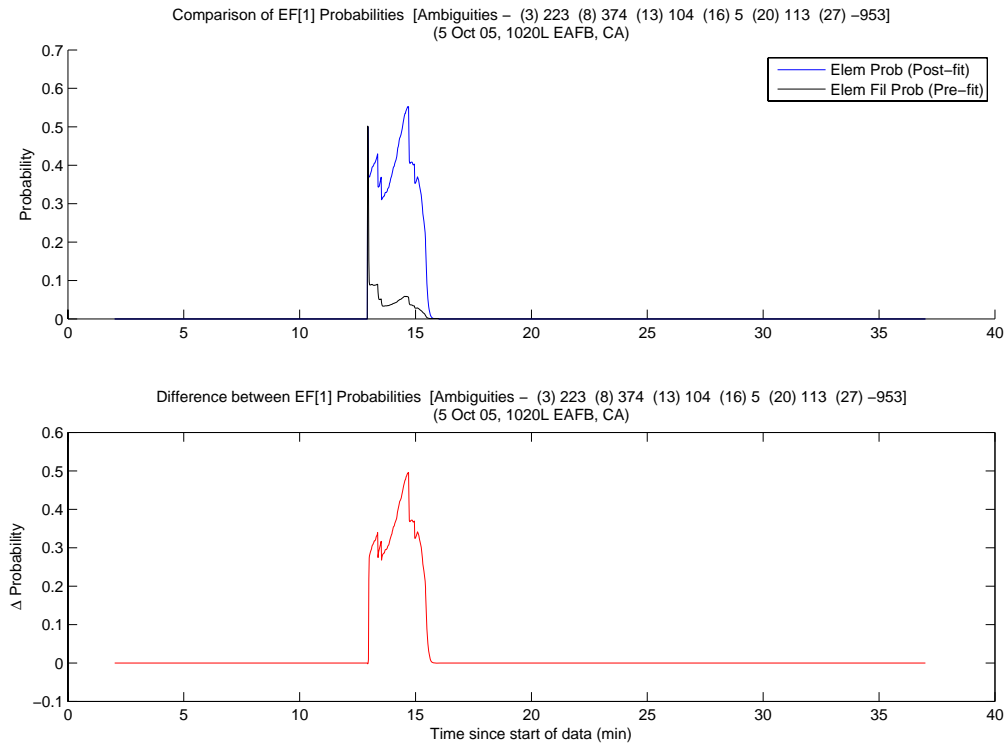


Figure A.166: Case F3.2: EF[1A] Probability Comparison

The MMAE position error is shown in Figure A.172. The MMAE solution started off relatively accurate, but diverged at 11 minutes when EF[7] absorbed over 50 percent of the probability. Next, the filter resets caused the MMAE position solution to remain inaccurate. Once the filter stabilized at approximately 18 minutes, the MMAE position solution remained accurate.

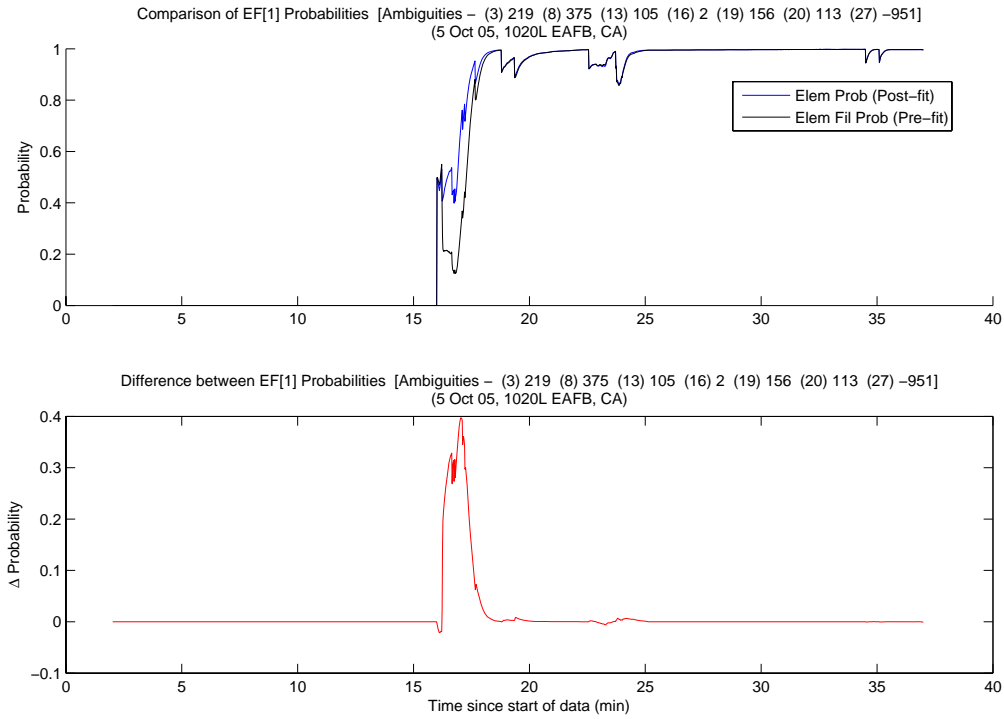


Figure A.167: Case F3.2: EF[1B] Probability Comparison

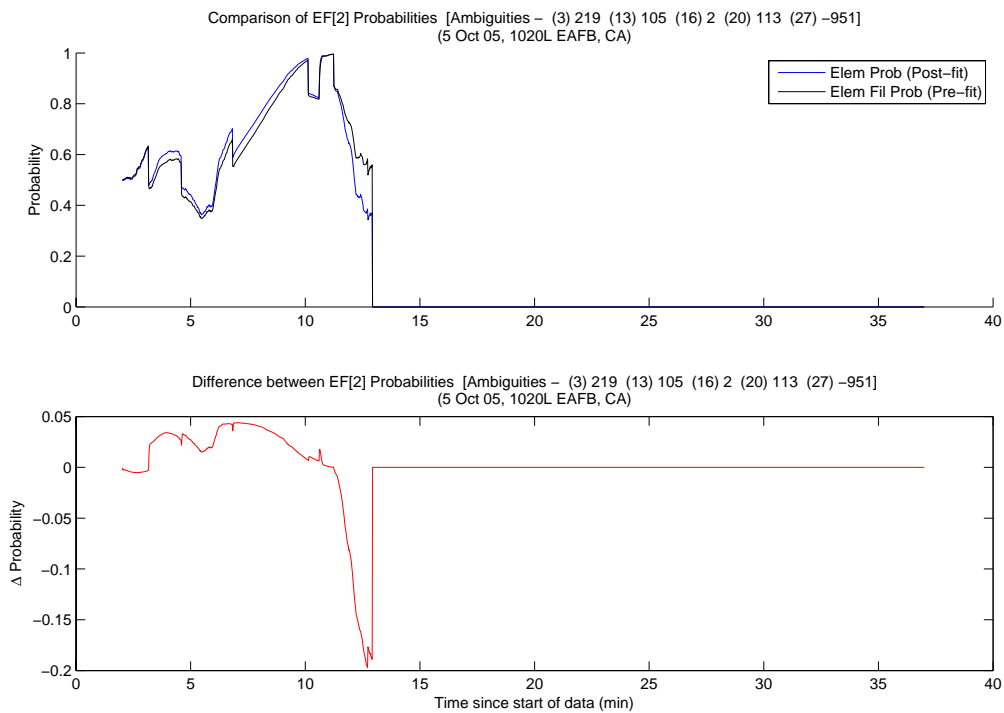


Figure A.168: Case F3.2: EF[2] Probability Comparison

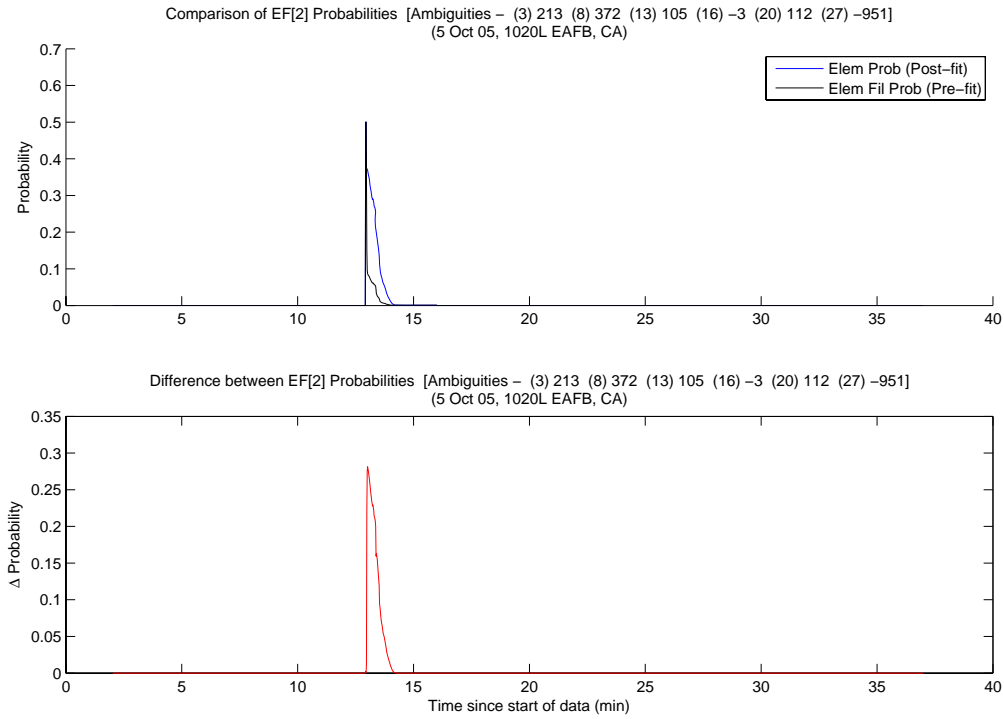


Figure A.169: Case F3.2: EF[2A] Probability Comparison

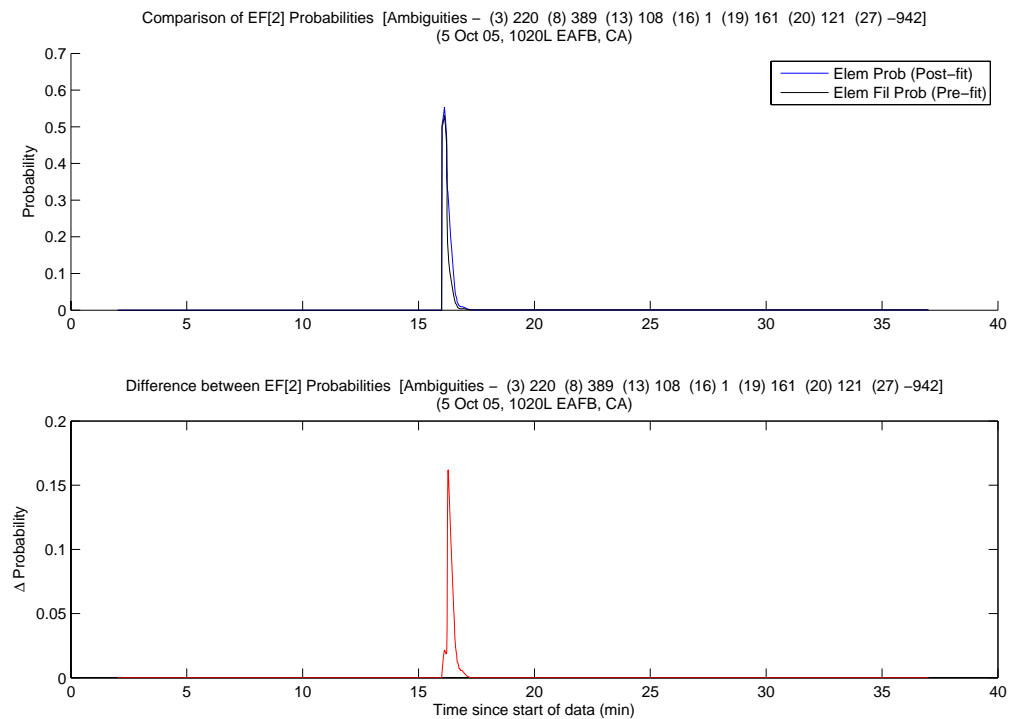


Figure A.170: Case F3.2: EF[2B] Probability Comparison

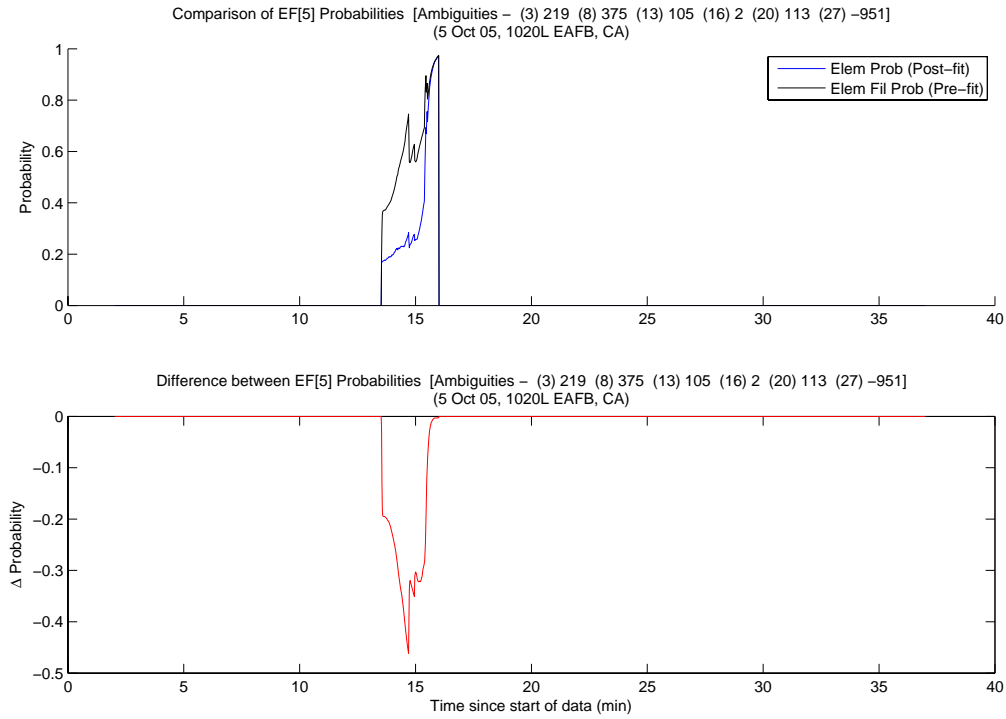


Figure A.171: Case F3.2: EF[5] Probability Comparison

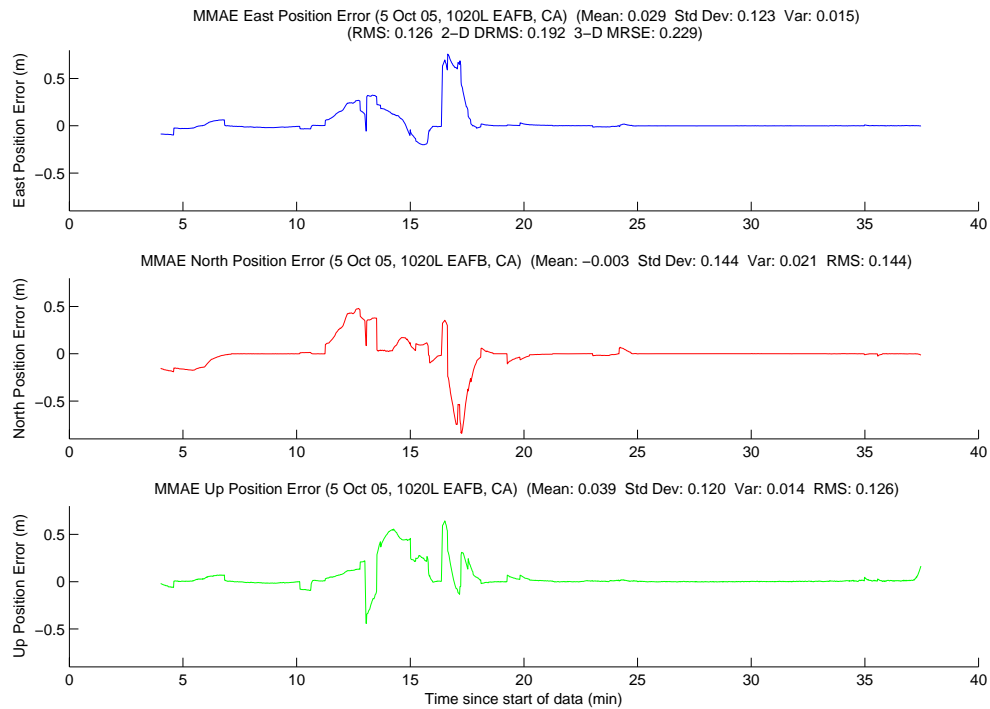


Figure A.172: Case F3.2: MMAE Position Error

Finally, Figure A.173 displays the position error plot for the minimum indicator variable. Unfortunately, since the minimum indicator incorrectly chose an erroneous ambiguity set after the first reset, the MRSE was outside the centimeter-level requirement. However, if the datalink had functioned properly, the minimum indicator would have remained “locked” to the true ambiguity set and met the centimeter requirement. The only jump in the minimum indicator plot is at 15.5 minutes when it chose the wrong set for 40 seconds. There was no minimum indicator from 13 to 15.5 minutes. Lastly, Table A.10 gives a summary of each method’s position error.

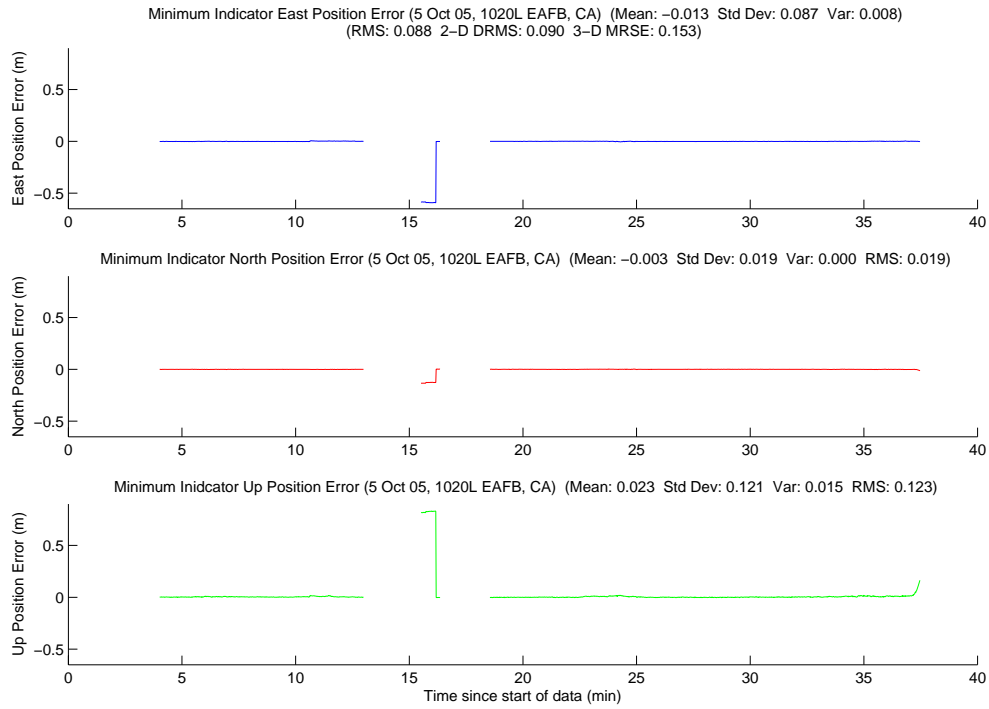


Figure A.173: Case F3.2: Minimum Indicator Position Error

Table A.10: Case F3.2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.133	0.180	0.073	0.121	0.083	0.164	0.265	0.322
MMAE	0.029	0.123	-0.003	0.144	0.039	0.120	0.192	0.229
Min. Ind.	-0.013	0.087	-0.003	0.019	0.023	0.121	0.090	0.153

A.2.7 Case F4.1: Flight Test 4. Case F4.1 was flown on 6 October 2005. The flight terminated early due to problems unrelated to the test hardware or software described in this thesis. Therefore, there was only one data run collected and it was 38 minutes long. The data run began at 1112L and the aircraft ranged from 7.5 to 3648.9 meters apart. Figure A.174 shows the minimum indicator North, East, Down relative positions.

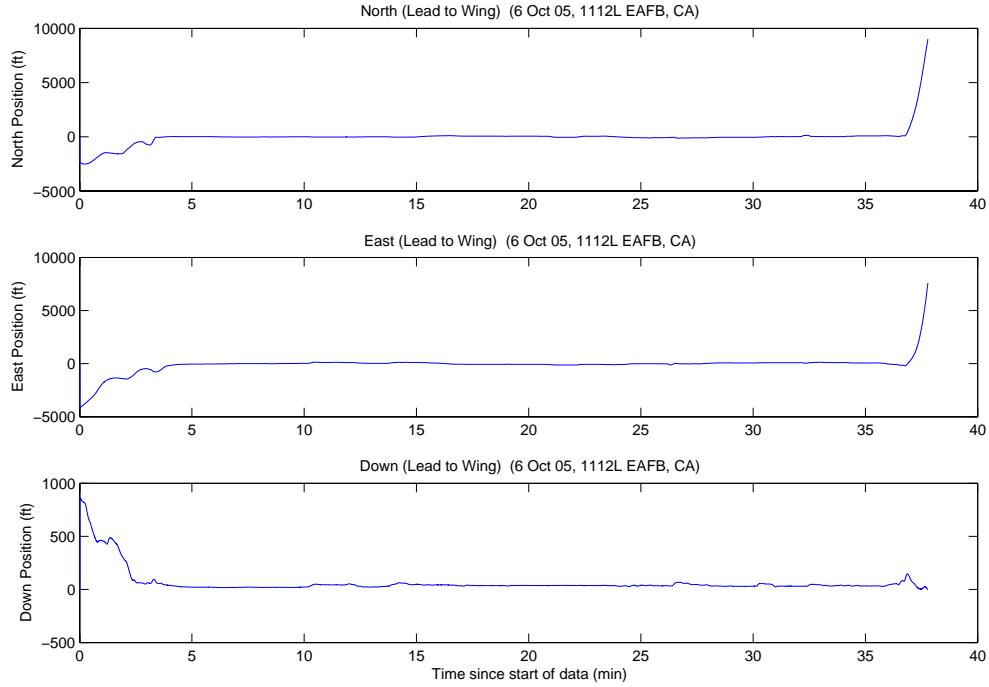


Figure A.174: Case F4.1: North, East, Down Relative Positions

The LAMBDA function returned a total of 31 unique ambiguity sets during the data run. The system was started once airborne when the aircraft were 1530 meters apart. Eight satellites were visible when the data run started. The true ambiguity set was the third set returned from the LAMBDA function, 47 seconds after the first ambiguity set was returned. The minimum indicator initially chose an incorrect ambiguity set at 4 minutes. It changed to the true ambiguity set 50 seconds later, and remained “locked” for the remainder of the data run. The true set was returned from LAMBDA as number one, 2120 of 2168 epochs (97.8 percent), with a maximum ratio of 37.7. The MMAE selected the true ambiguity set at three minutes, a full two minutes before the minimum indicator. When the minimum indicator selected the incorrect set at four minutes, the true set had already absorbed 92 percent of the conditional probability in the MMAE.

Figures A.175, A.176, and A.177 show the floating filter errors. The floating filter’s 1σ estimate of the error is outside the plot scales in all three figures.

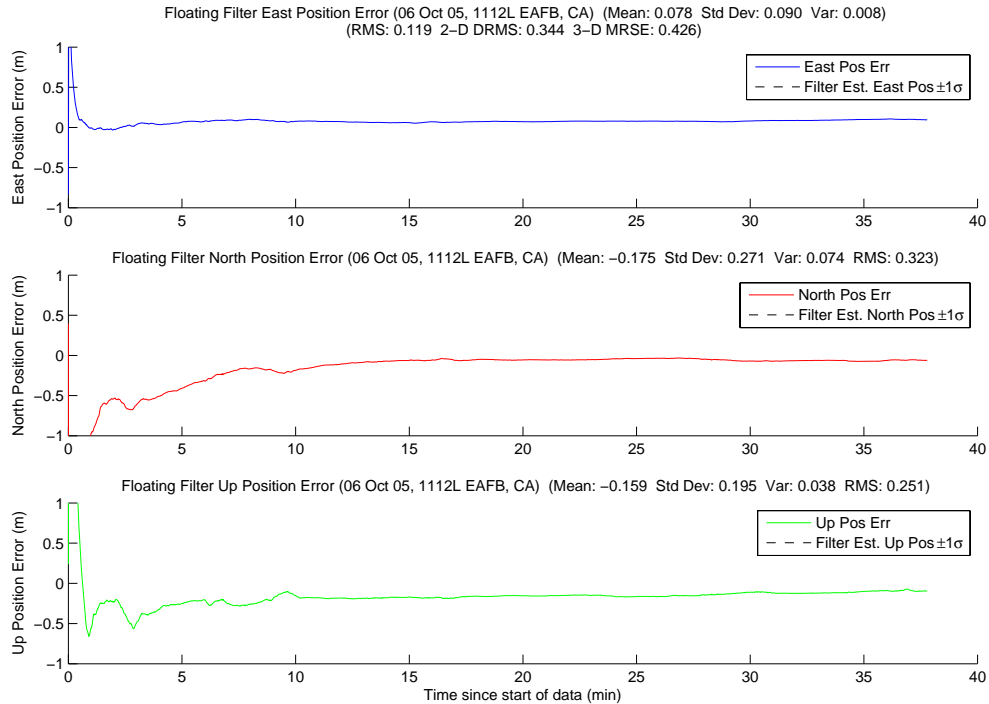


Figure A.175: Case F4.1: Floating Filter Position Errors

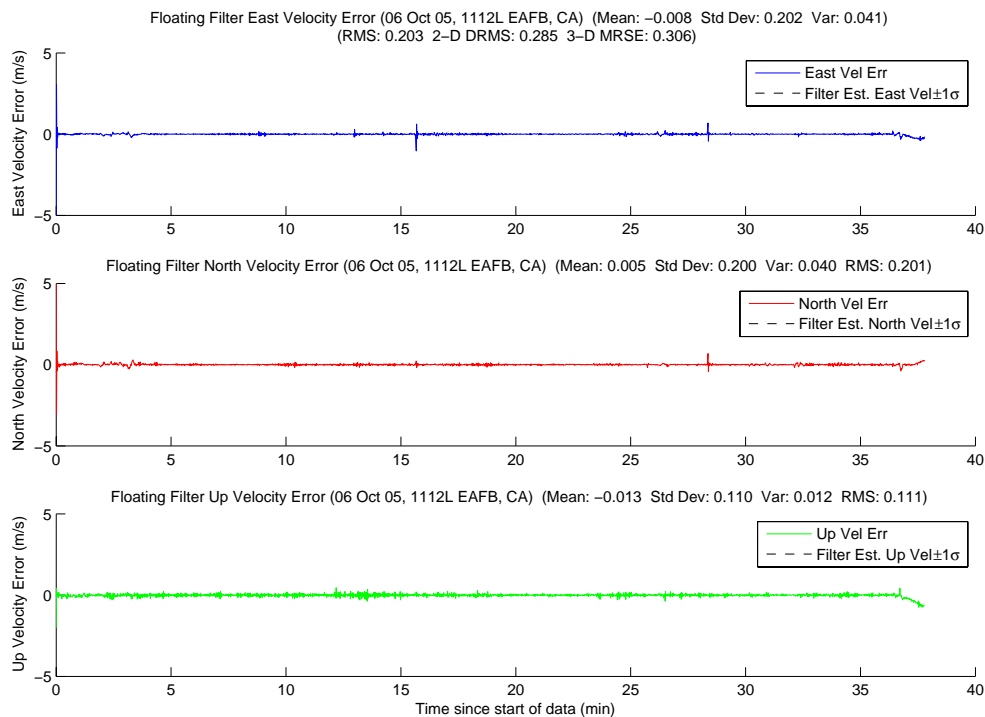


Figure A.176: Case F4.1: Floating Filter Velocity Errors

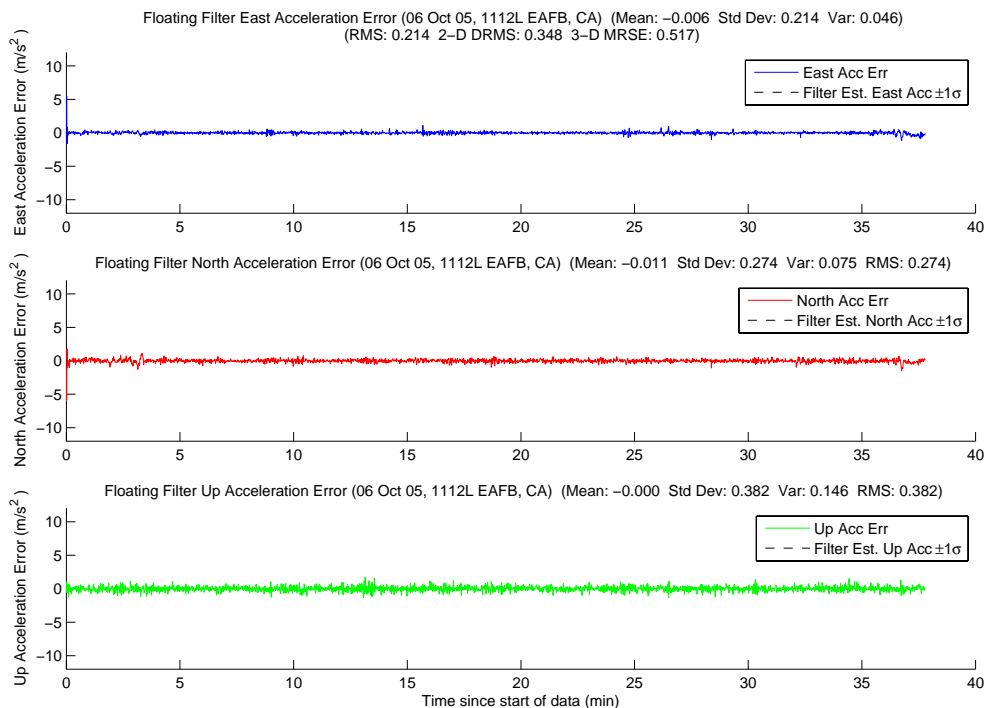


Figure A.177: Case F4.1: Floating Filter Acceleration Errors

Figure A.178 displays PRN 16 measurement information and the floating filter ambiguity estimate. There is no information for PRN 16 at 6.5 minutes for approximately 20 seconds since the satellite was masked. The rest of the satellite's plots follow in Figures A.179, A.180, A.181, A.182, A.183, A.184, A.185, and A.186.

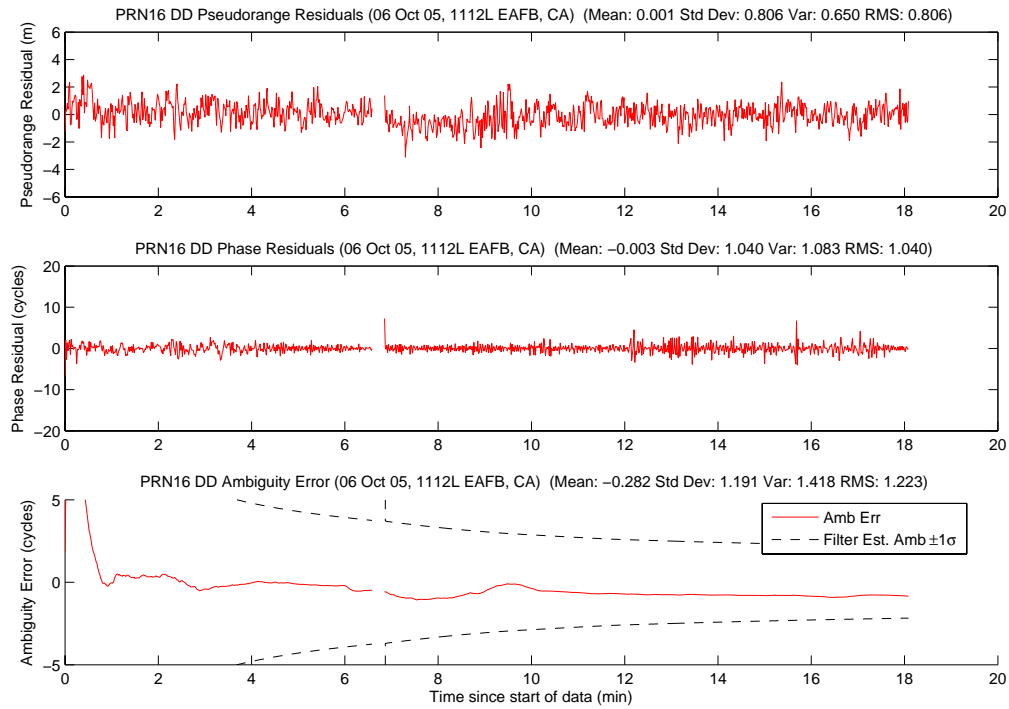


Figure A.178: Case F4.1: Satellite 16 Measurements

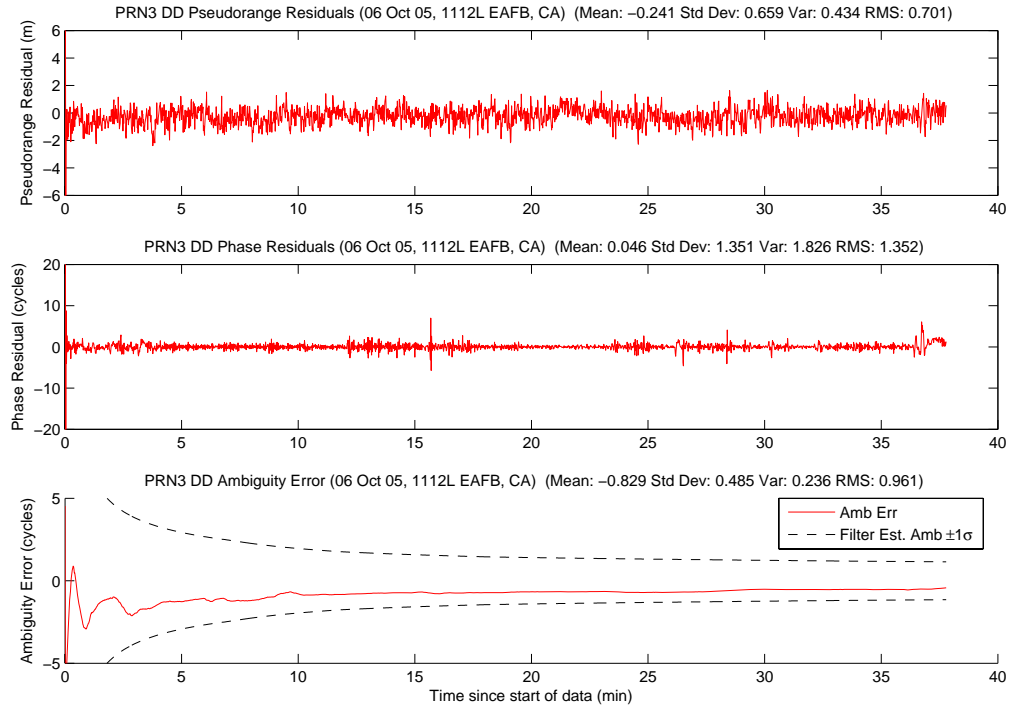


Figure A.179: Case F4.1: Satellite 3 Measurements

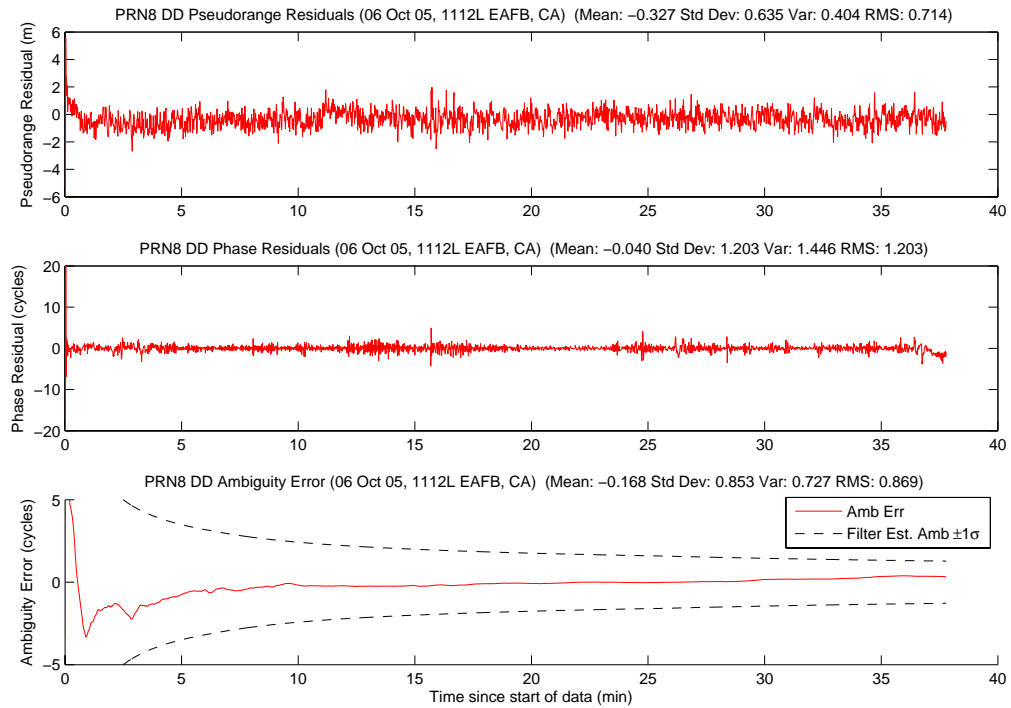


Figure A.180: Case F4.1: Satellite 8 Measurements

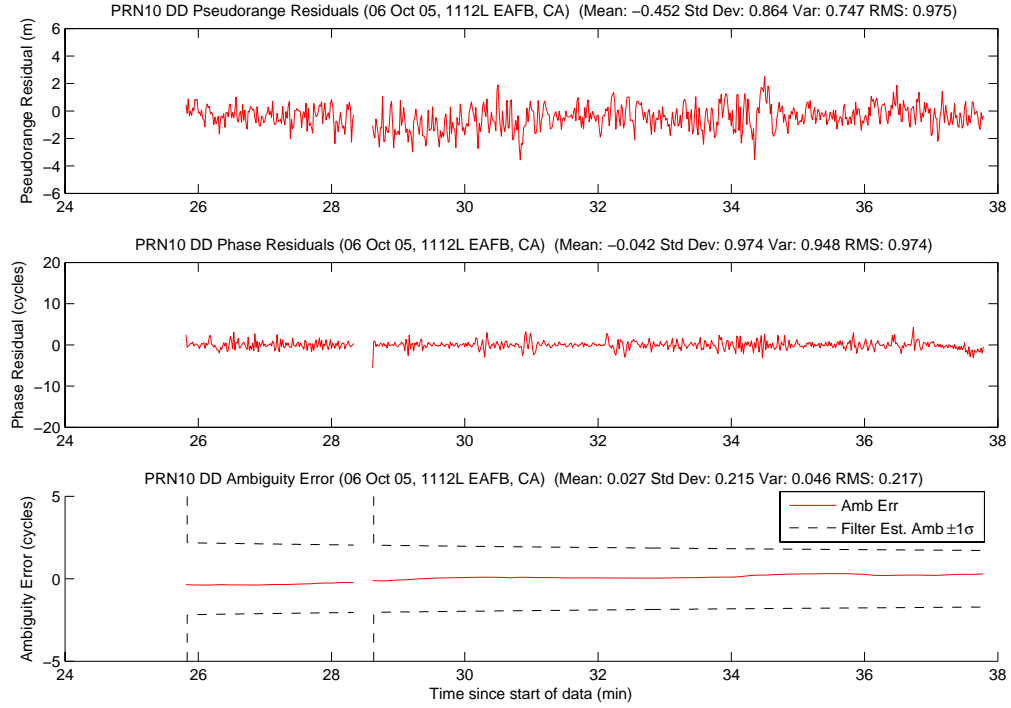


Figure A.181: Case F4.1: Satellite 10 Measurements

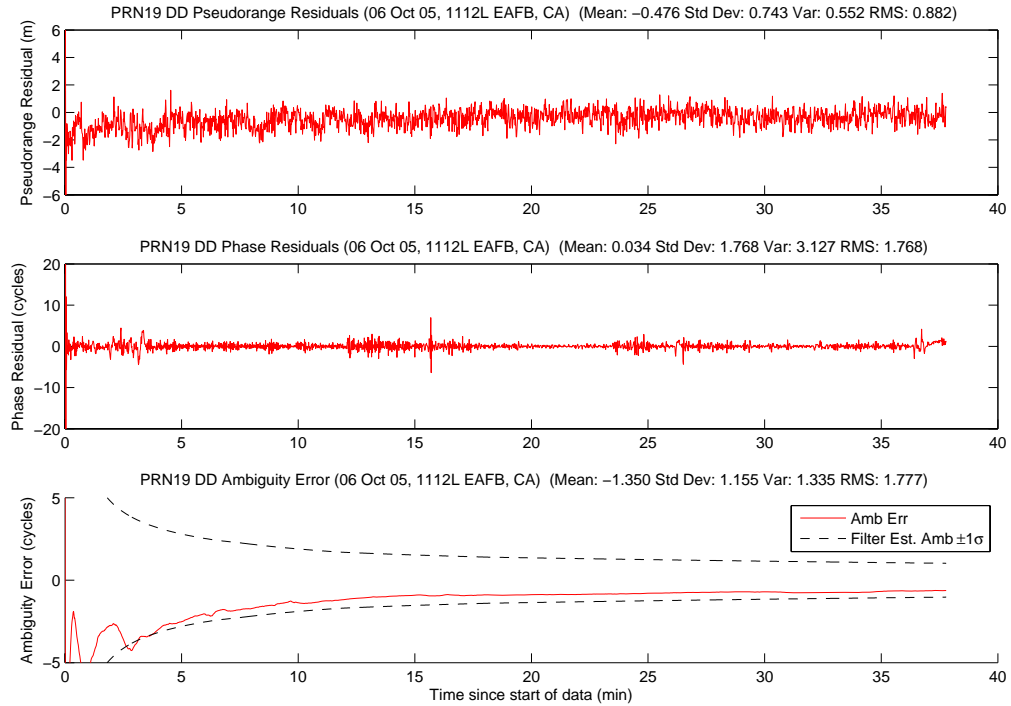


Figure A.182: Case F4.1: Satellite 19 Measurements

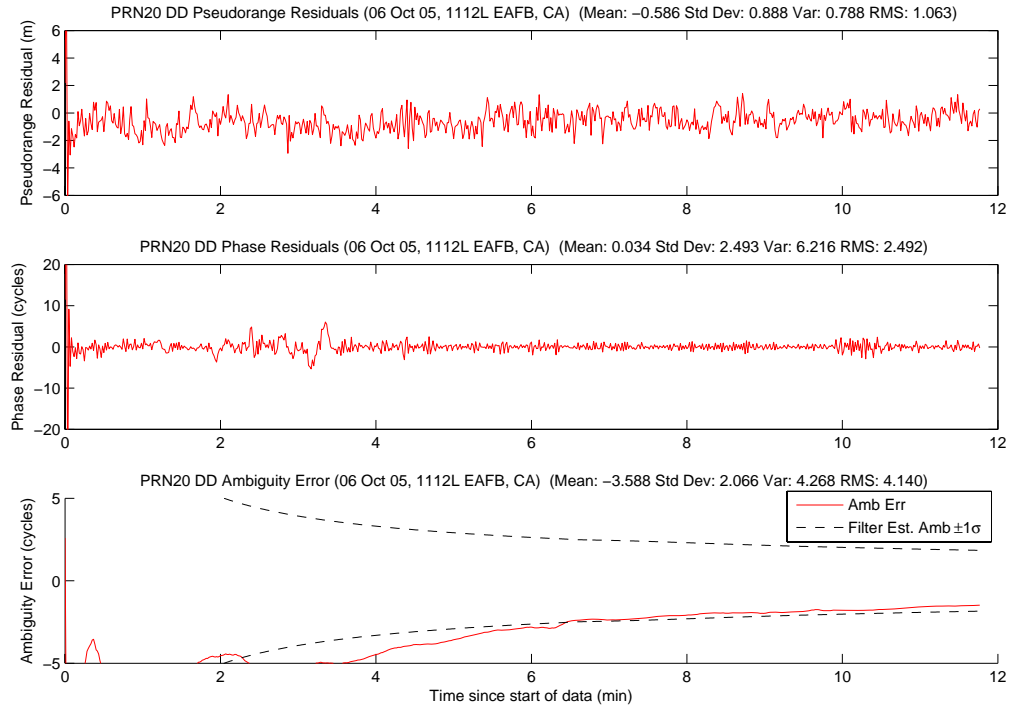


Figure A.183: Case F4.1: Satellite 20 Measurements

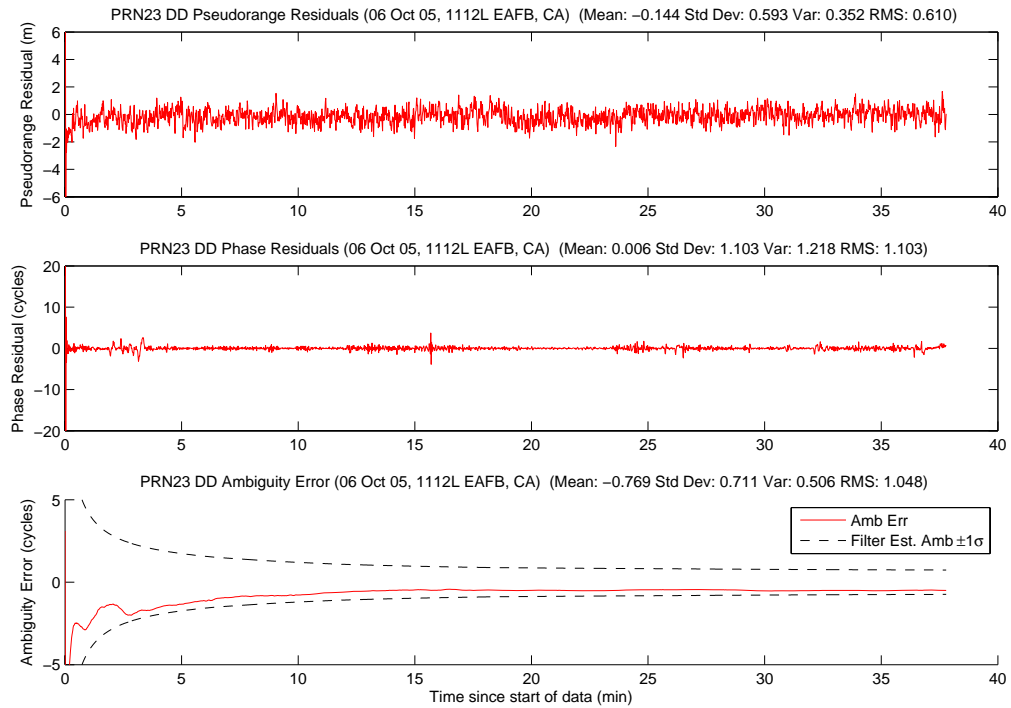


Figure A.184: Case F4.1: Satellite 23 Measurements

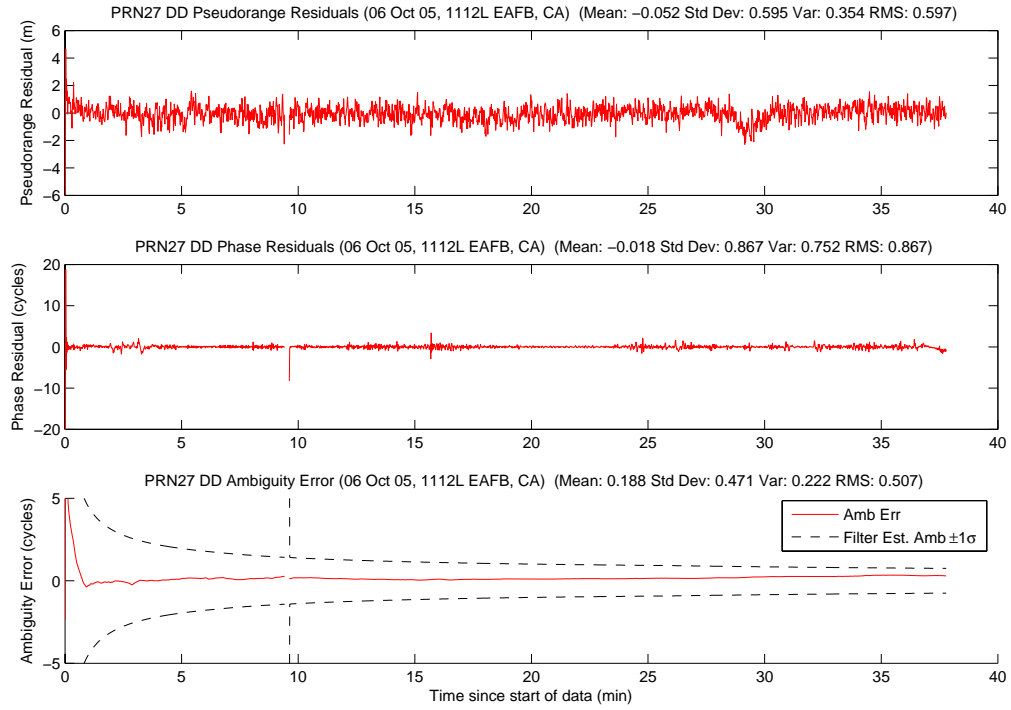


Figure A.185: Case F4.1: Satellite 27 Measurements

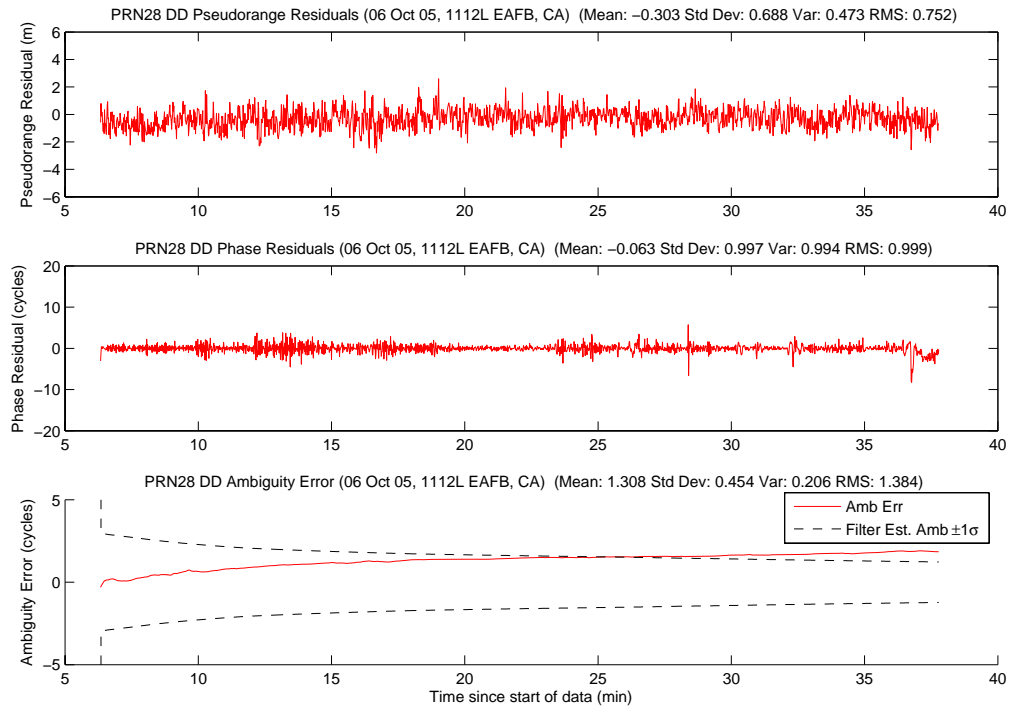


Figure A.186: Case F4.1: Satellite 28 Measurements

Figure A.187 shows the number of visible satellites. The satellite dropouts at 7, 10, and 28 minutes were from satellite masking. At 7 and 10 minutes, the aircraft were in close formation during simulated air refueling. The dropout at 28 minutes was PRN 10 which had risen two minutes earlier and was still low on the horizon.

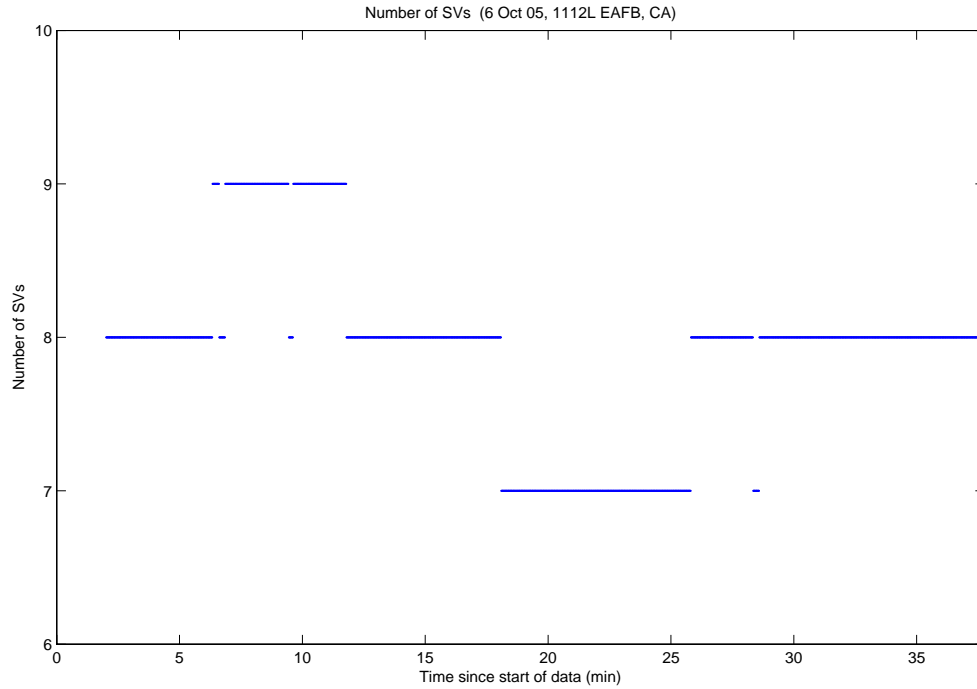


Figure A.187: Case F4.1: Visible Satellites

Figure A.188 shows all of the candidate ambiguity set's SOS residuals, and Figure A.189 displays the MMAE conditional probability. The probability plot shows a dip in the true ambiguity's probability at seven minutes. There was an erroneous set's residuals near the true set's residuals at that time in Figure A.188. Also, the dip coincided with PRN 16 masking. A similar event occurred between 26 and 28 minutes when PRN 10 rose above the elevation cutoff angle. This shows how a lower elevated satellite, which was affected more by atmospheric errors, can detrimentally affect the position solution. Finally, all of the residuals began to grow in the last 30 seconds of the data run when the aircraft were almost four kilometers apart.

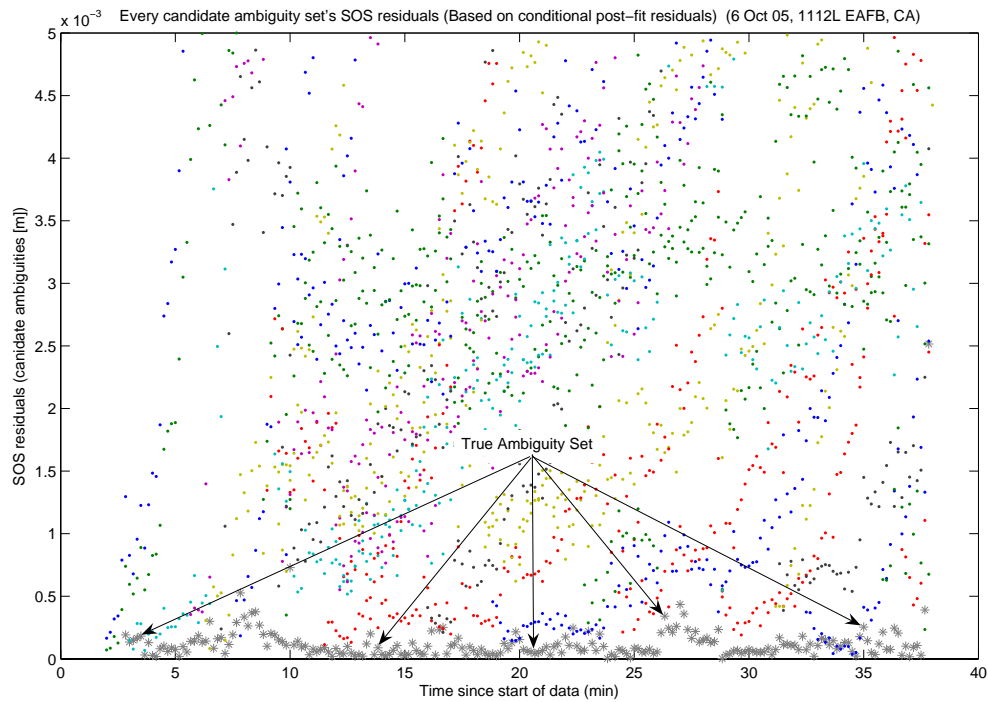


Figure A.188: Case F4.1: SOS Residuals

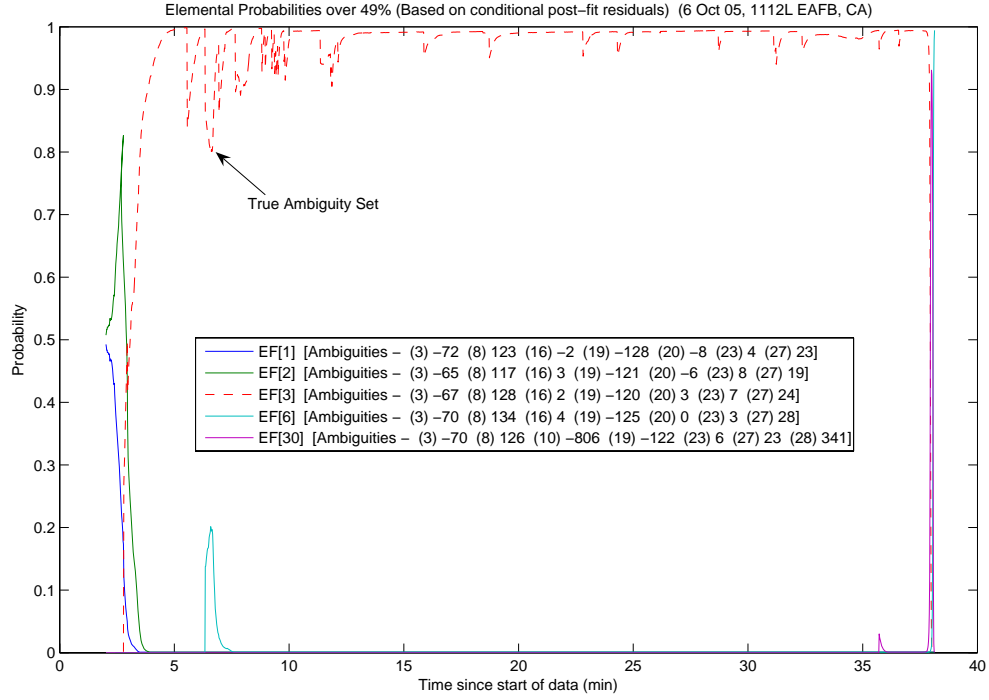


Figure A.189: Case F4.1: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.190 is EF[1], Figure A.191 is EF[2], and Figure A.192 is EF[3].

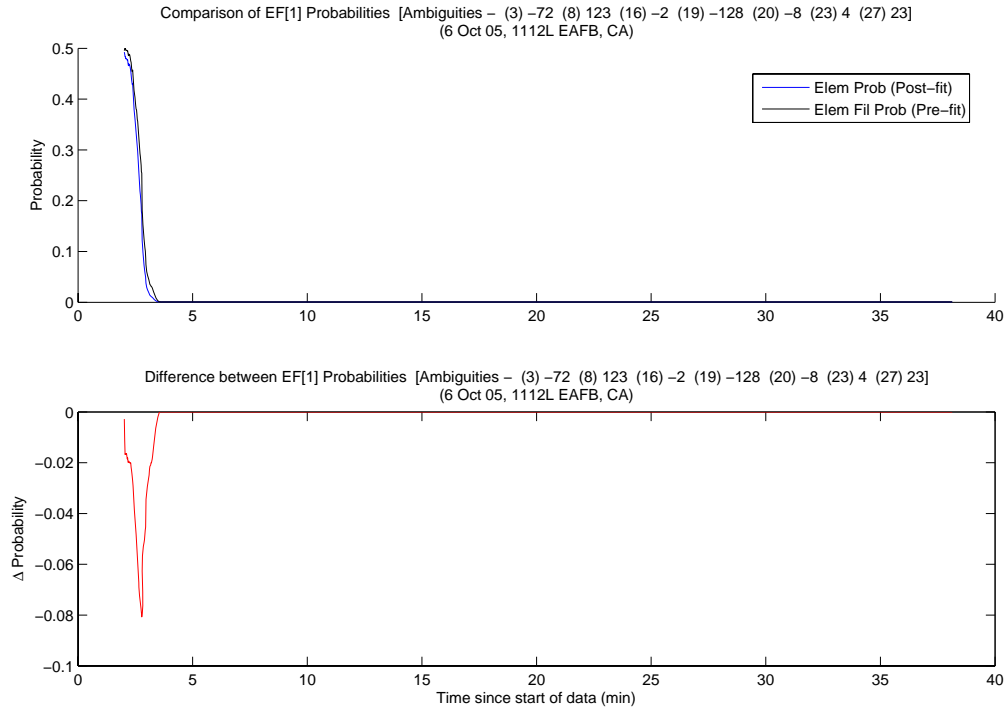


Figure A.190: Case F4.1: EF[1] Probability Comparison

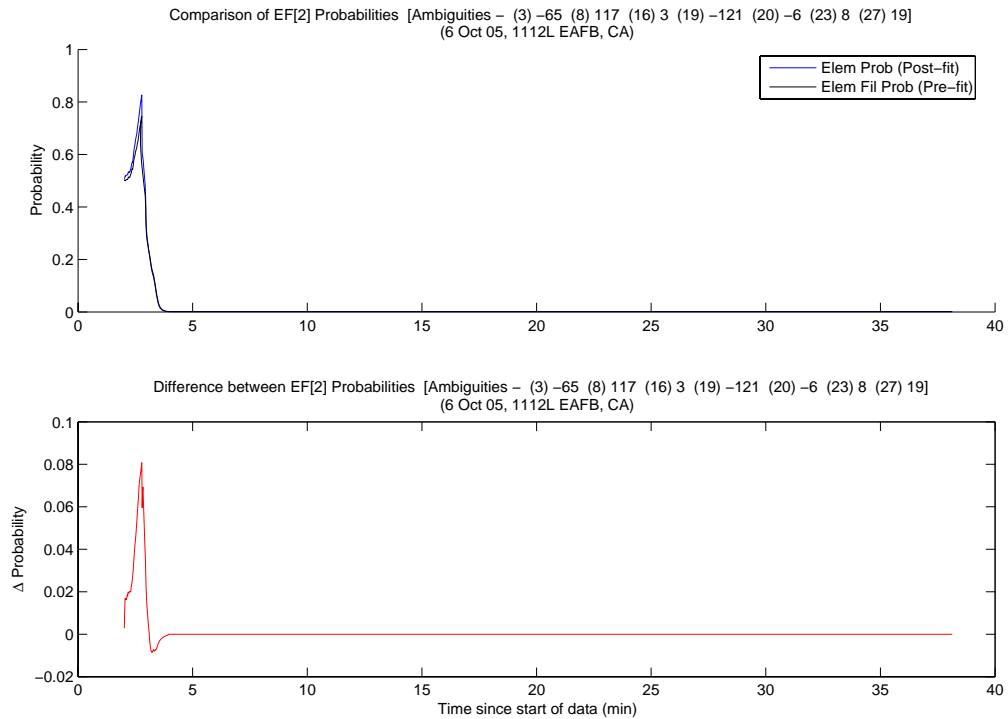


Figure A.191: Case F4.1: EF[2] Probability Comparison

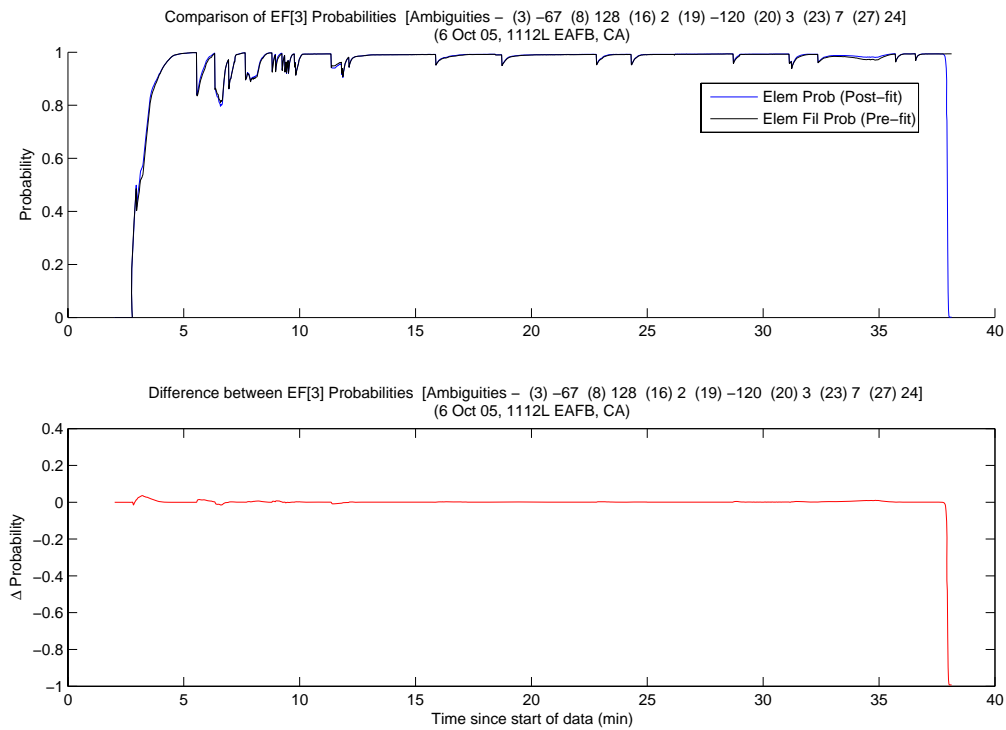


Figure A.192: Case F4.1: EF[3] Probability Comparison

The MMAE position error is shown in Figure A.193. The MMAE solution was accurate, except at five minutes when the true ambiguity set's probability dropped, and at the end of the data run when the aircraft began to separate. The last 30 second of data was included to show how the position errors grow with aircraft separation.

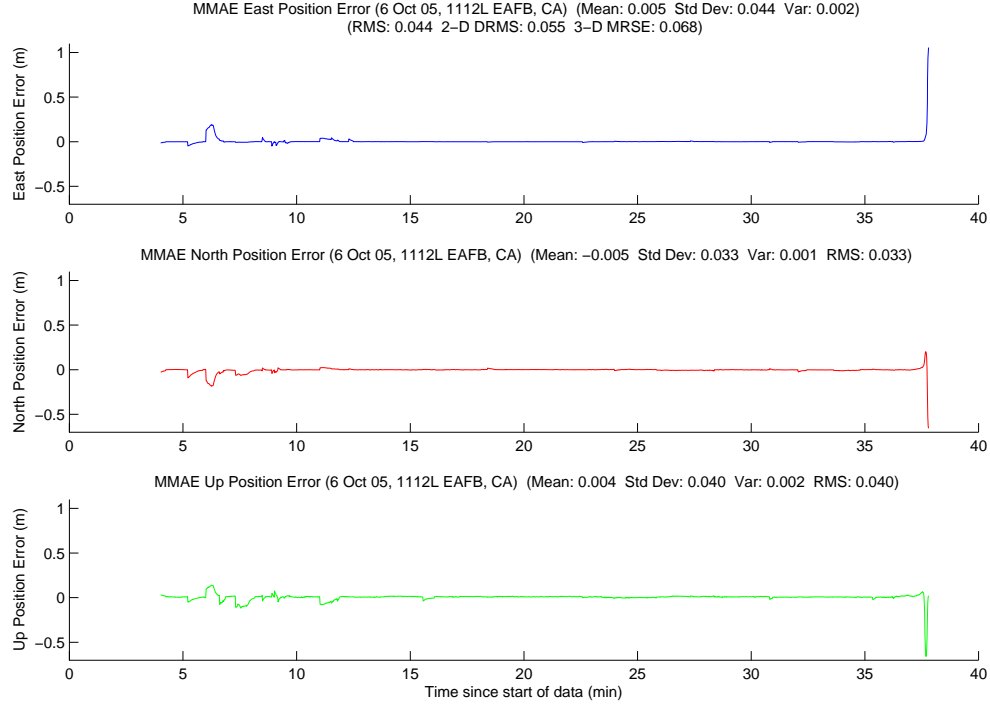


Figure A.193: Case F4.1: MMAE Position Error

Finally, Figure A.194 displays the position error plot for the minimum indicator variable. Unfortunately, since the minimum indicator chose an erroneous ambiguity at four minutes, the MRSE was outside the centimeter-level requirement. However, like Case 3.1, the filter was initialized when the aircraft were 1530 meters apart. Once the aircraft joined (the critical part of the flight) the MRSE was well within the centimeter requirement. It is also interesting to note that with the true ambiguity set, a position accuracy under 0.5 meters was possible when the aircraft were 3.5 kilometers apart. Lastly, Table A.11 gives a summary of each method's position error.

Table A.11: Case F4.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.078	0.090	-0.175	0.271	-0.159	0.195	0.344	0.426
MMAE	0.005	0.044	-0.005	0.033	0.004	0.040	0.055	0.068
Min. Ind.	-0.019	0.161	-0.009	0.075	-0.010	0.164	0.179	0.243

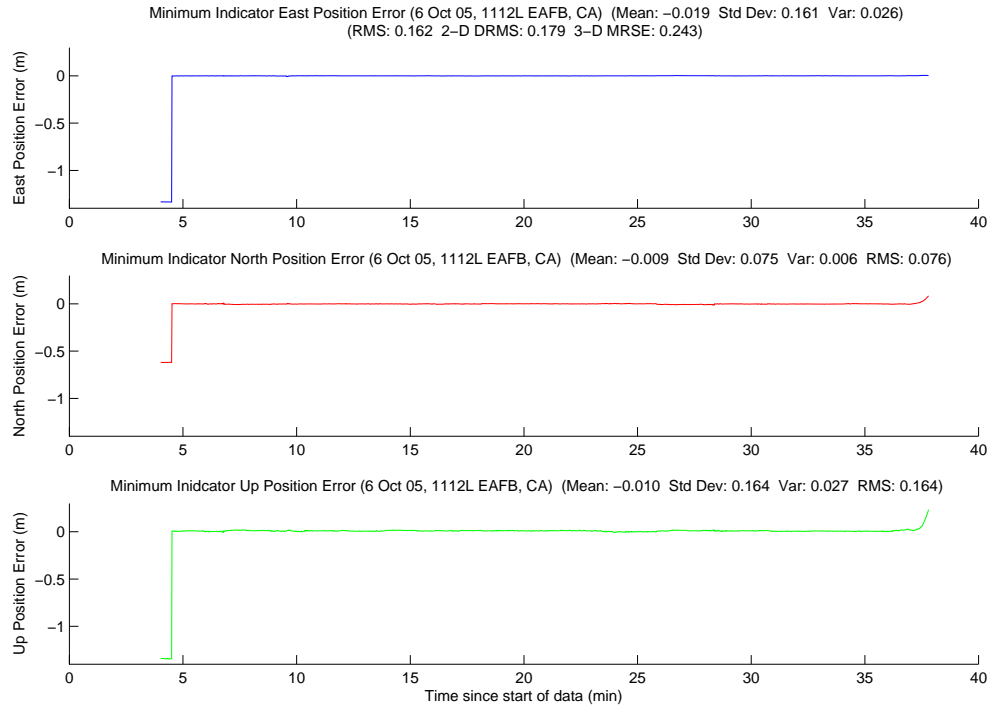


Figure A.194: Case F4.1: Minimum Indicator Position Error

A.2.8 Case F5.1: Flight Test 5, First Data Set . Case F5.1 was flown on 11 October 2005 at 1516L and lasted 59 minutes, 33 seconds. The weather during the test was clear with light winds. The aircraft ranged from 15.0 to 455.6 meters apart during the data run. Figure A.195 shows the minimum indicator North, East, Down relative positions.

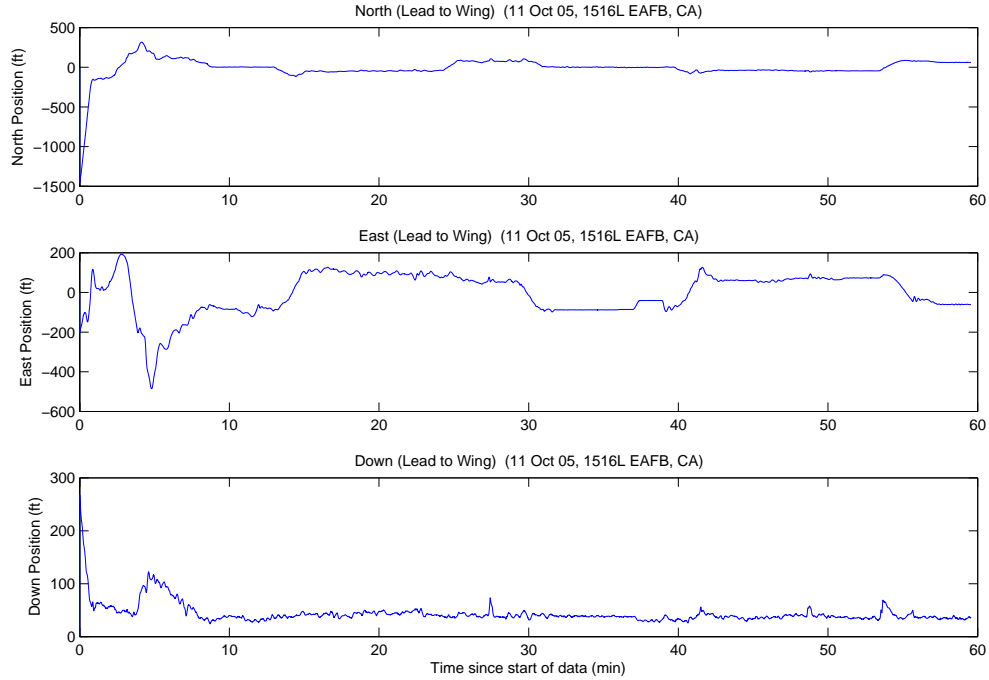


Figure A.195: Case F5.1: North, East, Down Relative Positions

The LAMBDA function returned a total of 18 unique ambiguity sets during the data run. The system was started once airborne when the aircraft were 456 meters apart, with ten visible satellites. During the middle of the data run, the system was tracking twelve visible satellites; the maximum the Javad receiver could track. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at 4 minutes. After the minimum indicator selected the true set it did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 3452 of 3452 epochs (100.0 percent), with a maximum ratio of 161.2.

Figures A.196, A.197, and A.198 show the floating filter errors. The floating filter performed extremely well, as it should with an average of eleven satellites. The floating filter's 1σ estimate of the error is outside the plot scales in all three figures.

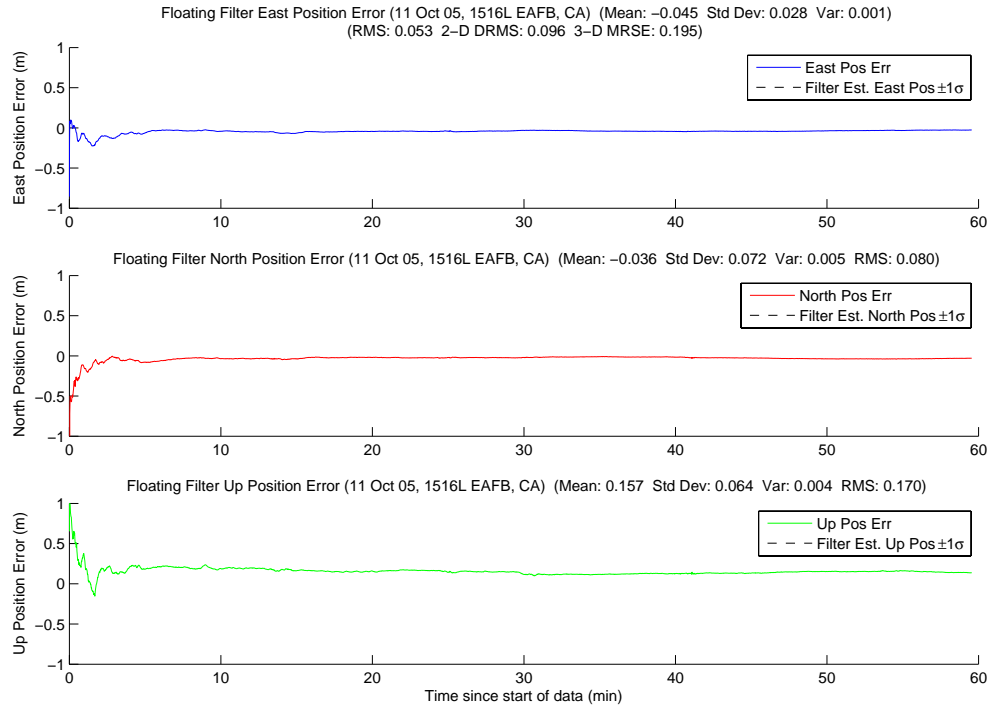


Figure A.196: Case F5.1: Floating Filter Position Errors

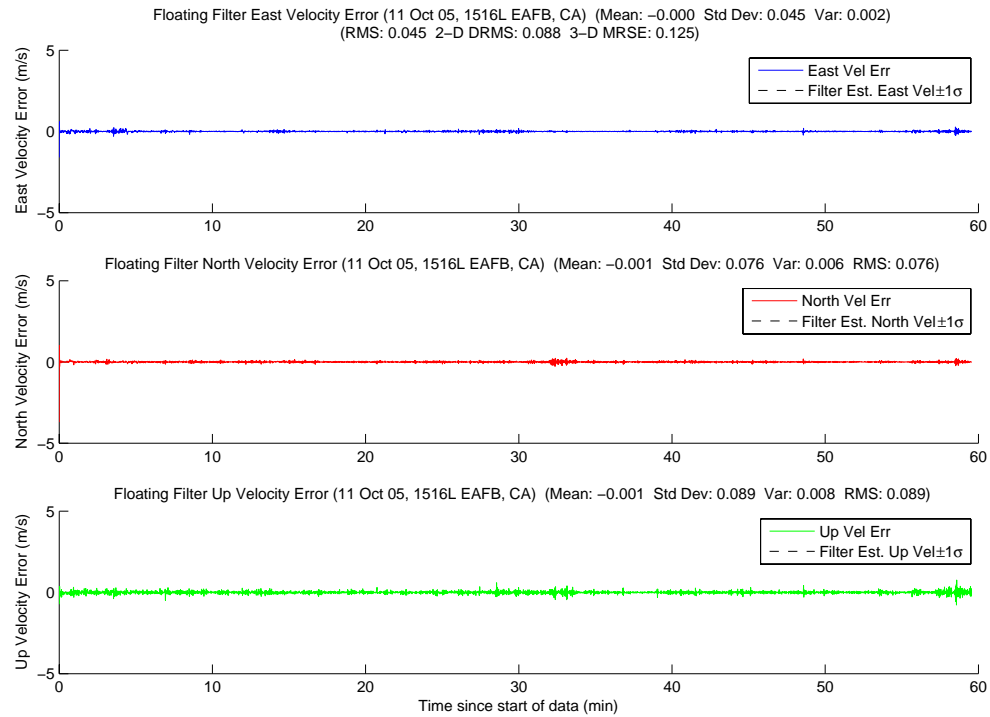


Figure A.197: Case F5.1: Floating Filter Velocity Errors

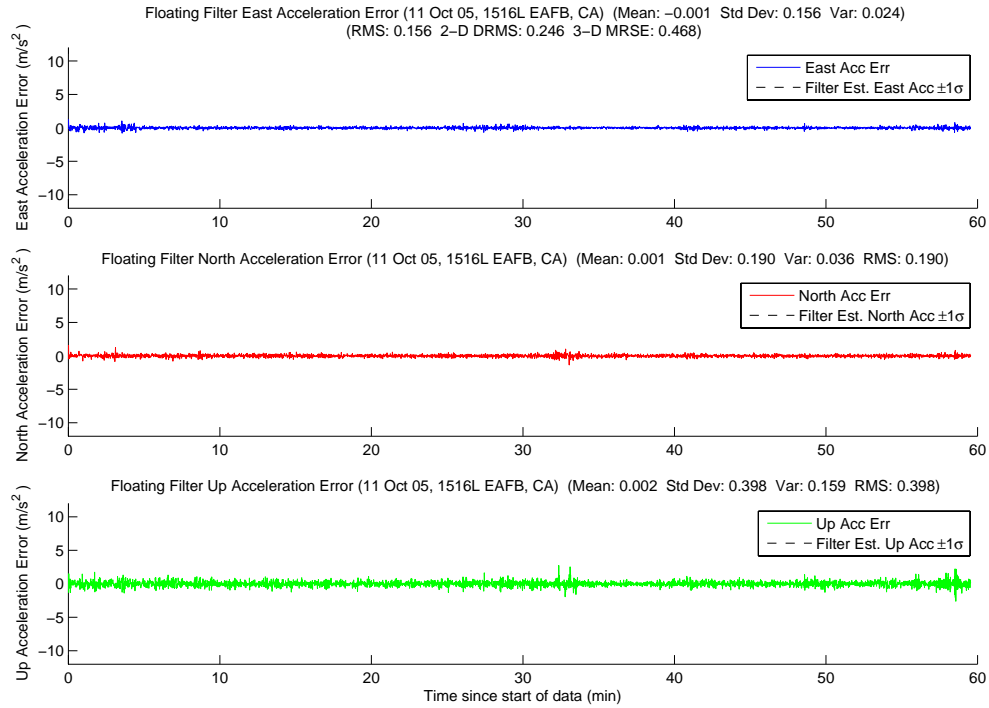


Figure A.198: Case F5.1: Floating Filter Acceleration Errors

Figure A.199 displays PRN 17 measurement information and the floating filter ambiguity estimate. The decrease in the floating filter 1σ ambiguity estimate at approximately 8.5 minutes, was due to a change in the maximum elevated satellite. The rest of the satellite's plots follow in Figures A.200, A.201, A.202, A.203, A.204, A.205, A.206, A.207, A.208, A.209, and A.210.

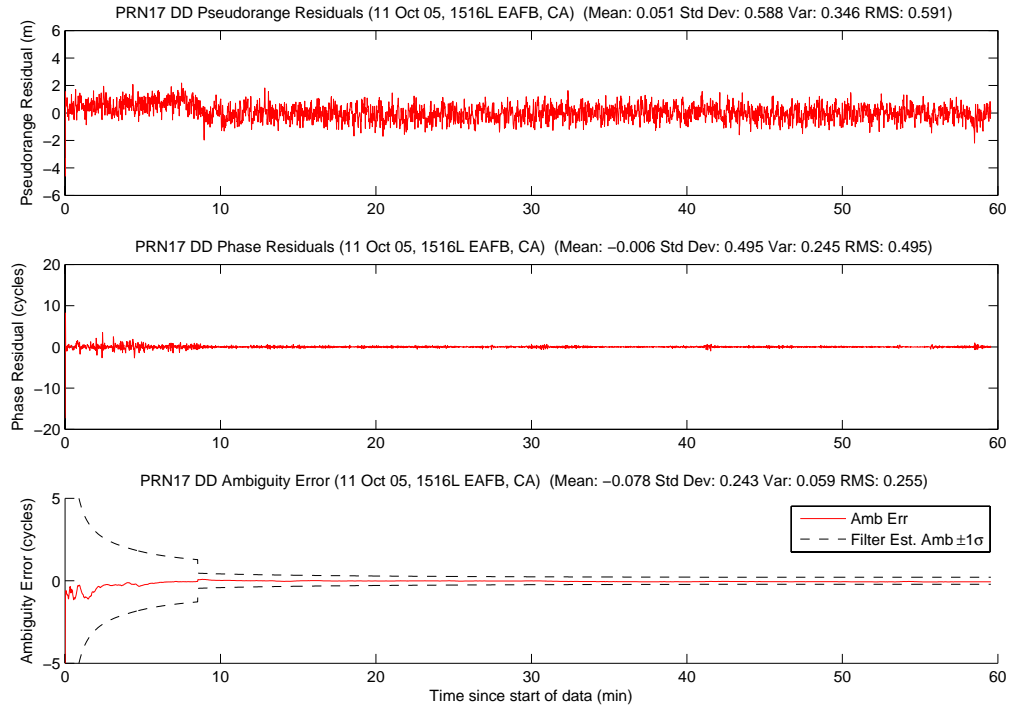


Figure A.199: Case F5.1: Satellite 17 Measurements

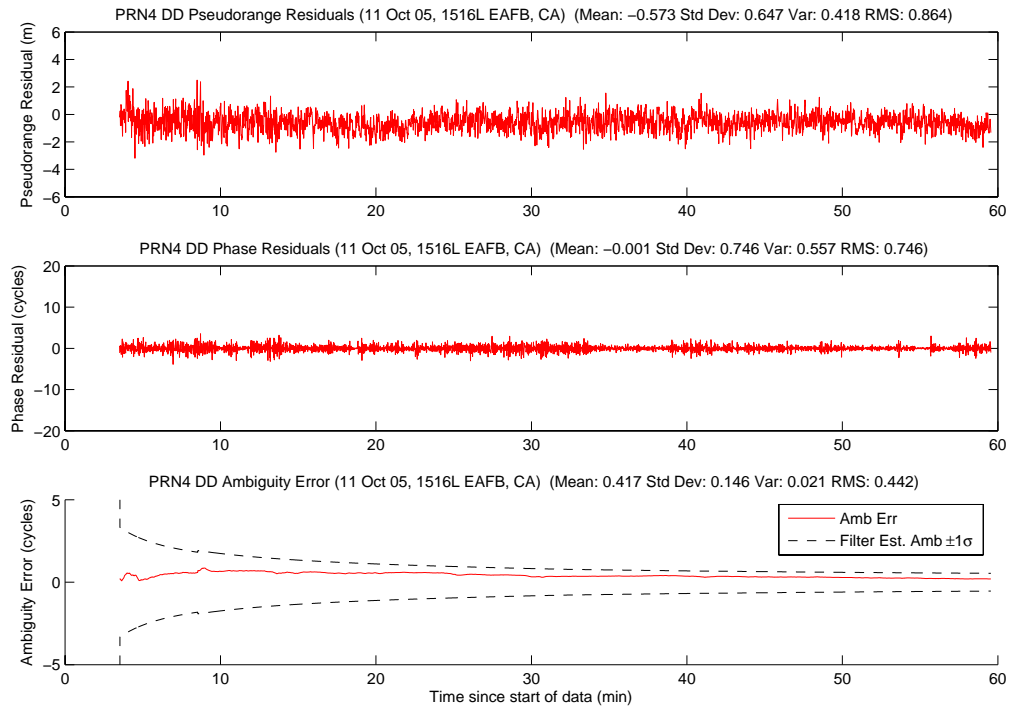


Figure A.200: Case F5.1: Satellite 4 Measurements

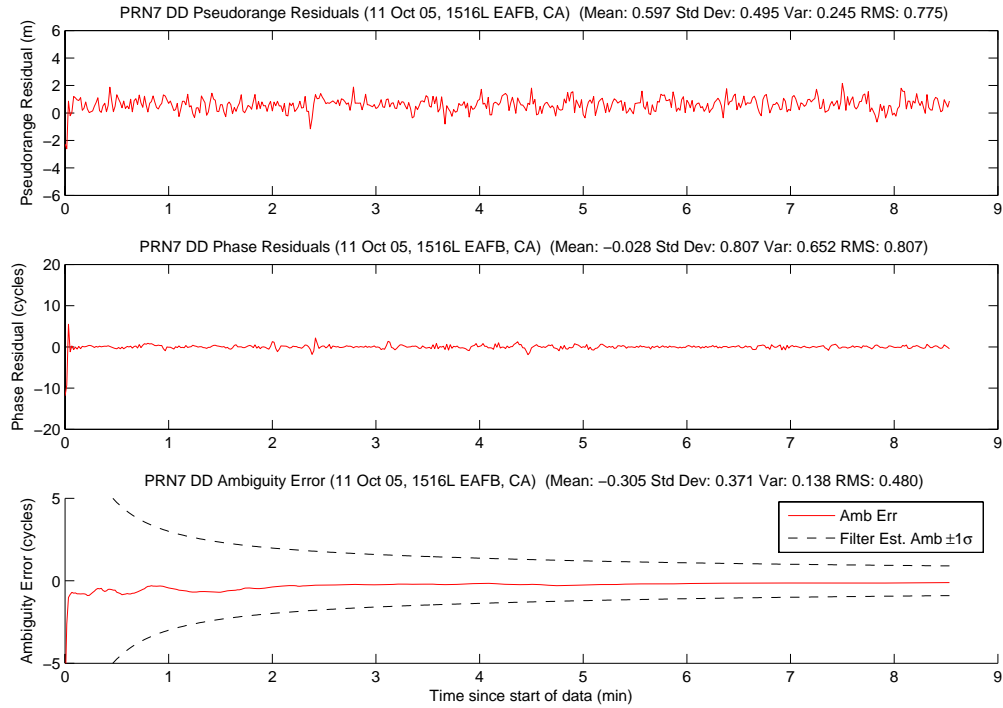


Figure A.201: Case F5.1: Satellite 7 Measurements

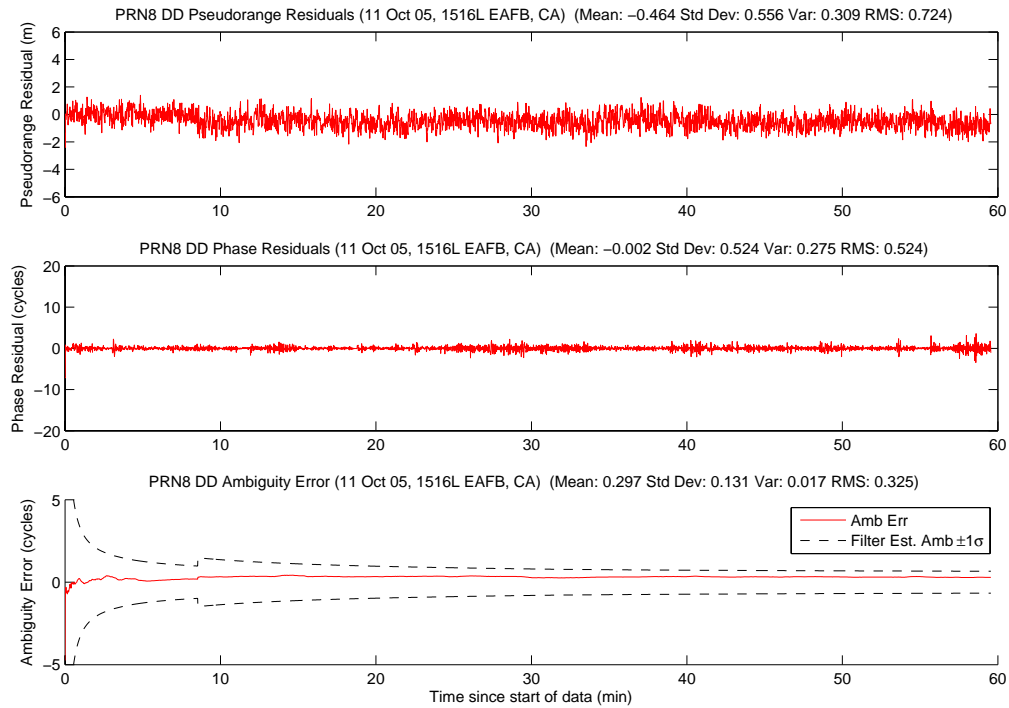


Figure A.202: Case F5.1: Satellite 8 Measurements

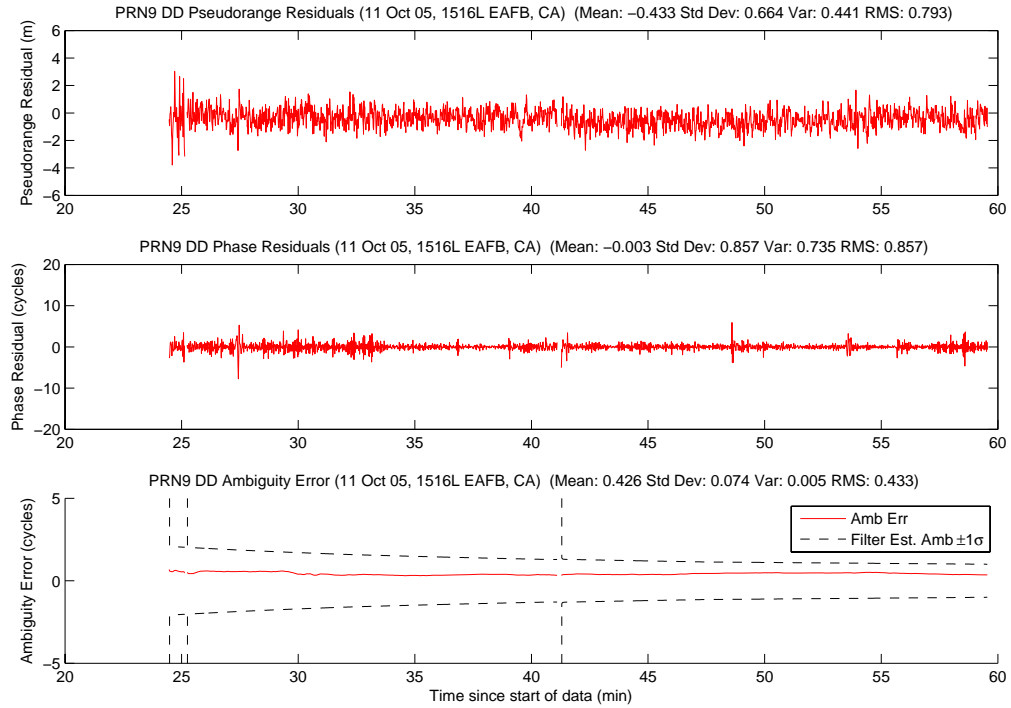


Figure A.203: Case F5.1: Satellite 9 Measurements

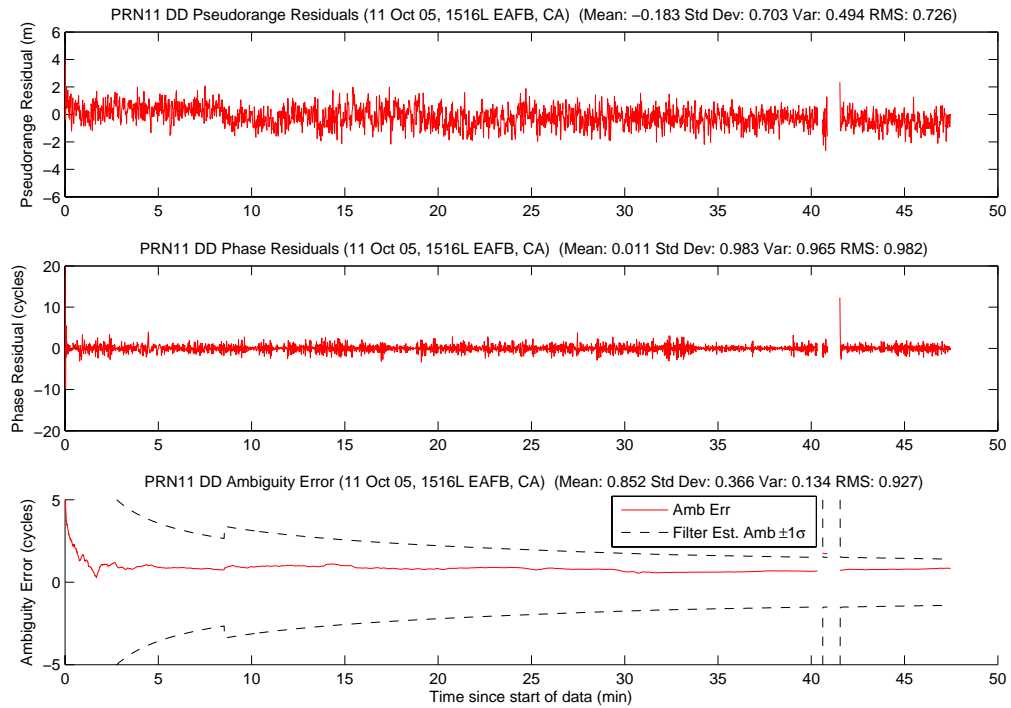


Figure A.204: Case F5.1: Satellite 11 Measurements

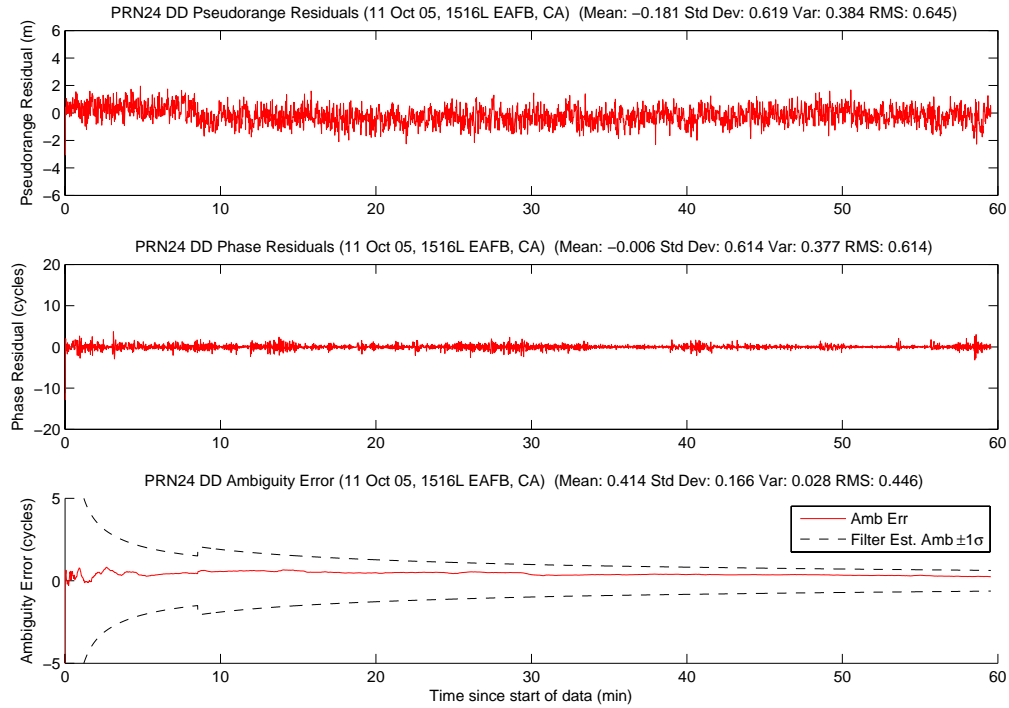


Figure A.205: Case F5.1: Satellite 24 Measurements

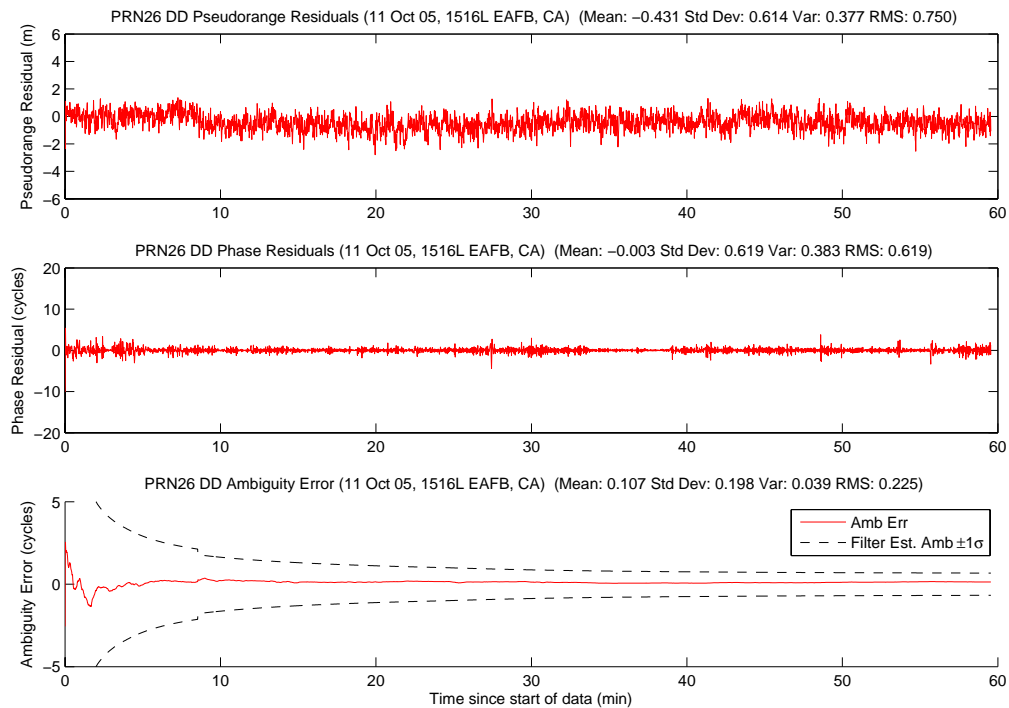


Figure A.206: Case F5.1: Satellite 26 Measurements

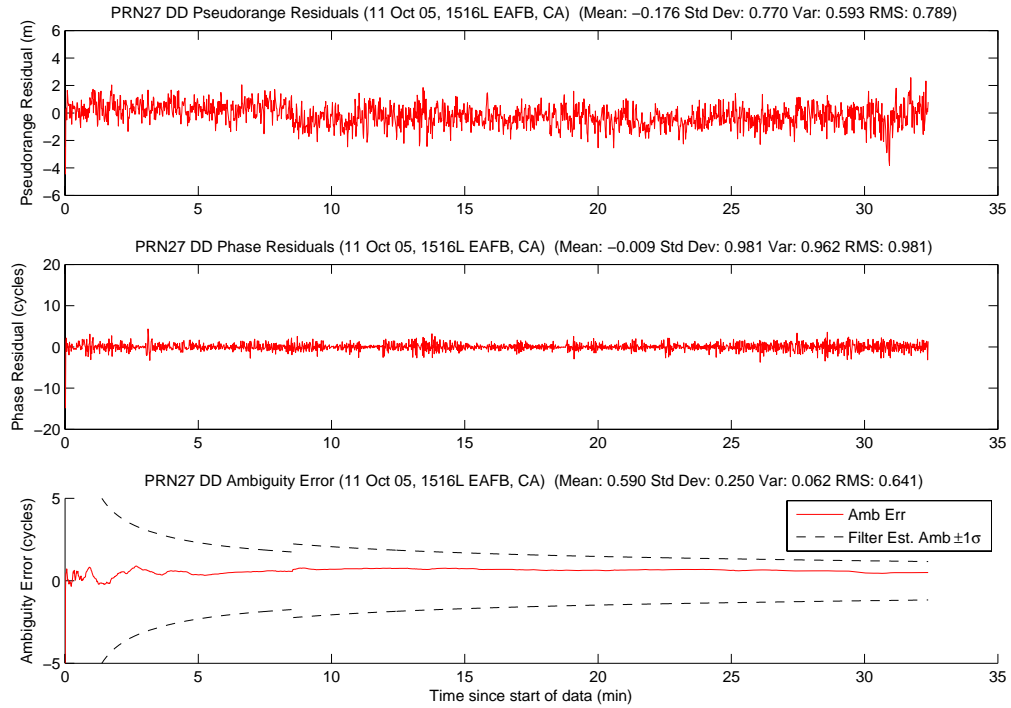


Figure A.207: Case F5.1: Satellite 27 Measurements

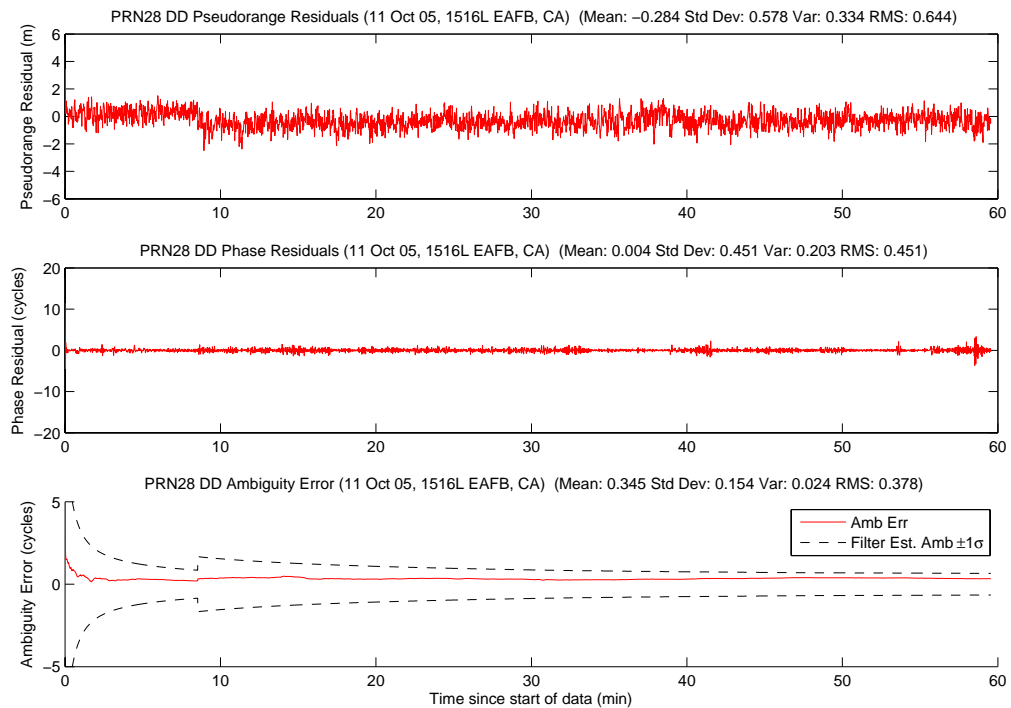


Figure A.208: Case F5.1: Satellite 28 Measurements

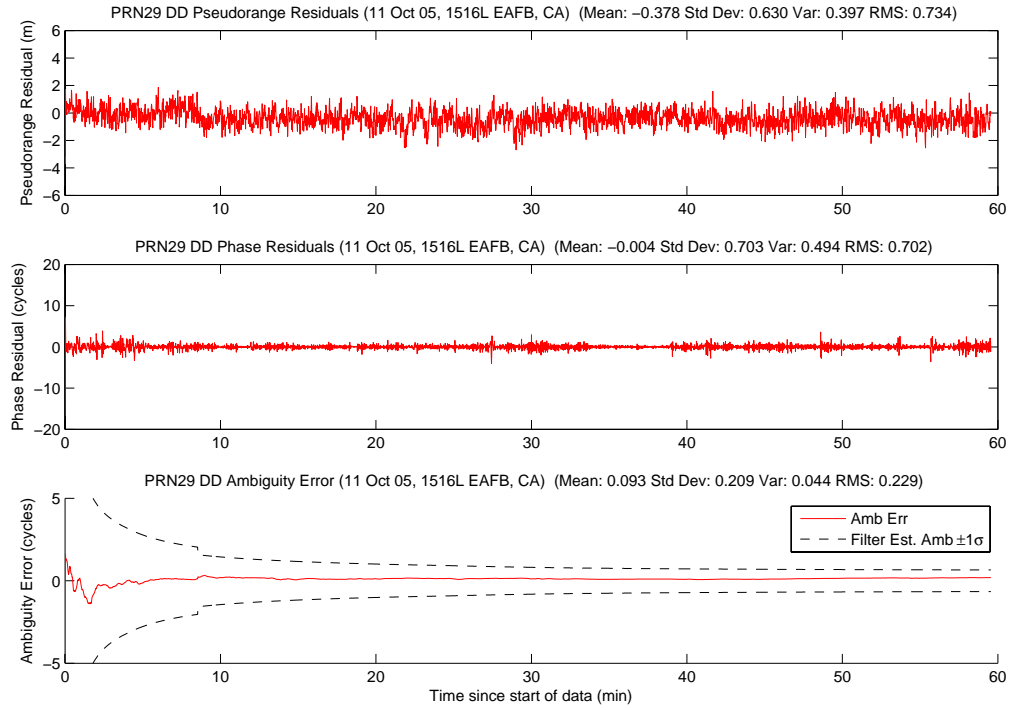


Figure A.209: Case F5.1: Satellite 29 Measurements

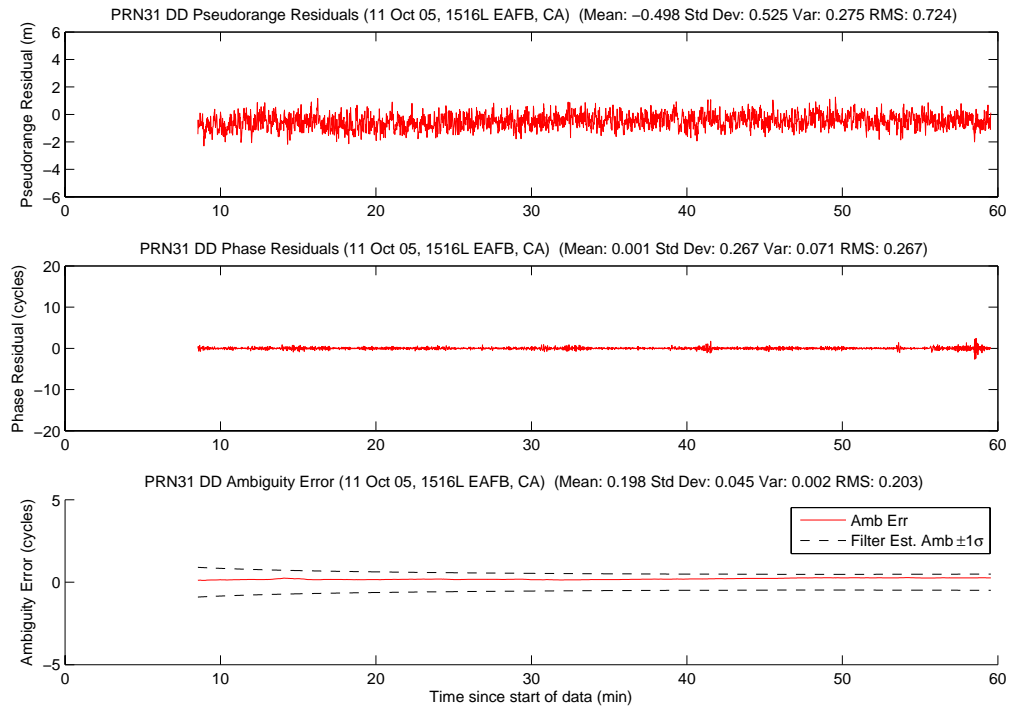


Figure A.210: Case F5.1: Satellite 31 Measurements

Figure A.211 shows the number of visible satellites. The satellite dropout at 25 minutes was due to an erroneously detected cycle slip on PRN 9, which had just risen one minute earlier and was still low on the horizon. The satellite dropouts at approximately 39 minutes were due to masking of lower elevated satellites, PRN 9 and 11. One of the dropouts, at 40.8 minutes, was due to a cycle slip on PRN 11, seven minutes prior to the satellite setting below ten degrees.

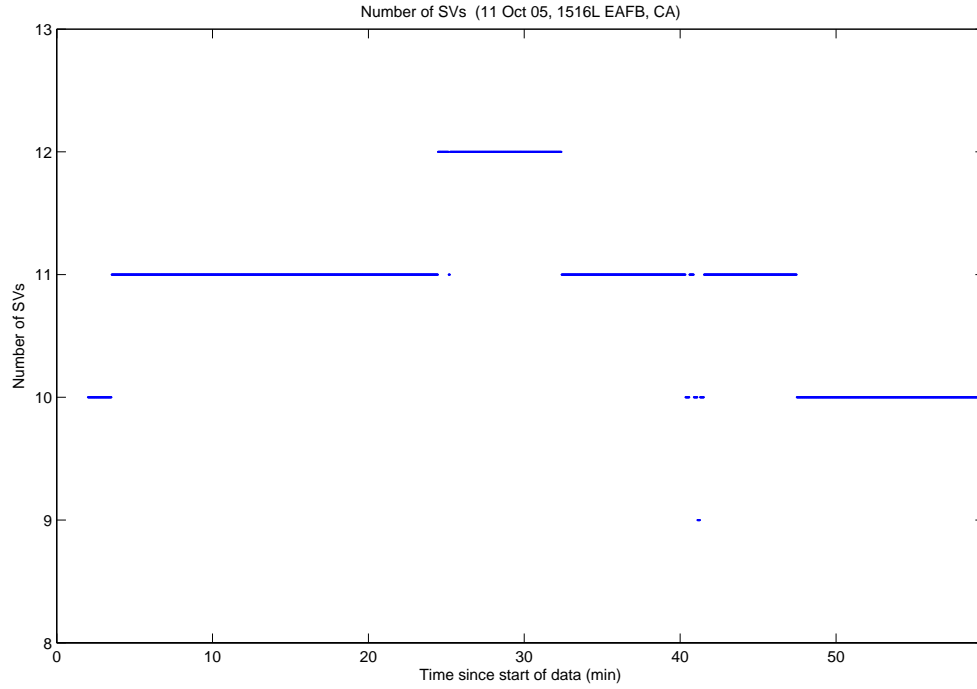


Figure A.211: Case F5.1: Visible Satellites

Figure A.212 shows all of the candidate ambiguity set's SOS residuals. The importance of the number of satellites is dramatically evident in the residual plot. Even from the beginning of the data run, there was not a single erroneous set's residuals near the true set. As the number of satellites increase, it becomes much more difficult for an incorrect ambiguity set's measurements to “fit” the position solution. Starting at approximately 25 minutes into the data run, when there were 12 satellites visible, no incorrect set's residuals are visible in the plot. Figure A.212 is similar to widelane SOS residuals plot.

Figure A.213 displays the MMAE conditional probability. The true ambiguity set maintained over 49 percent of the total probability for the entire data run. After the first two dips, the true set's probability did not drop below 90 percent.

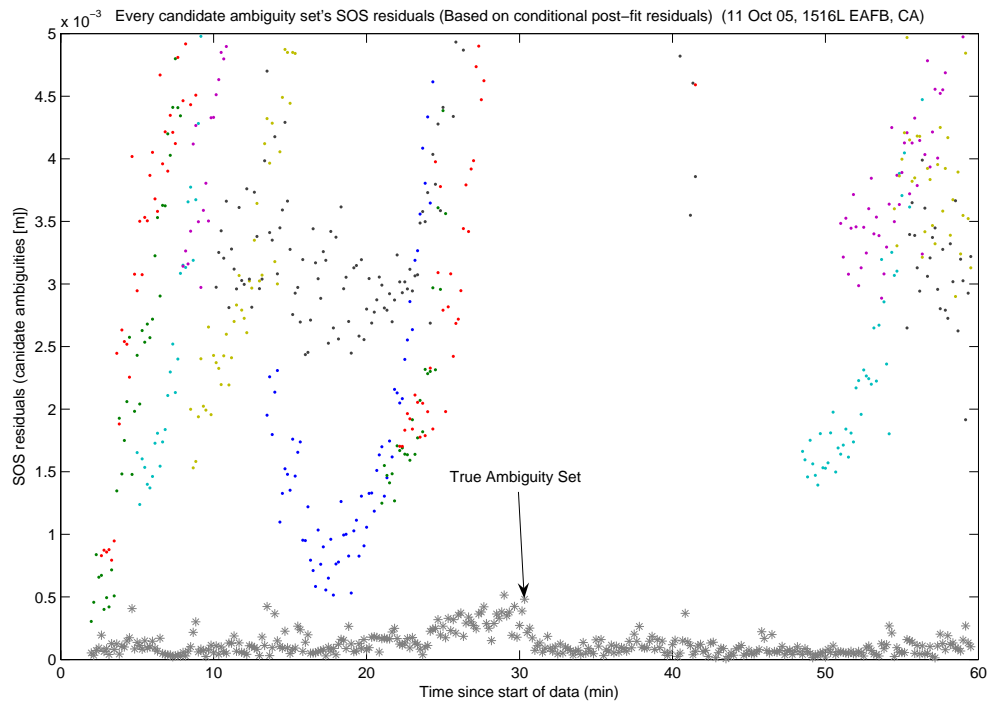


Figure A.212: Case F5.1: SOS Residuals

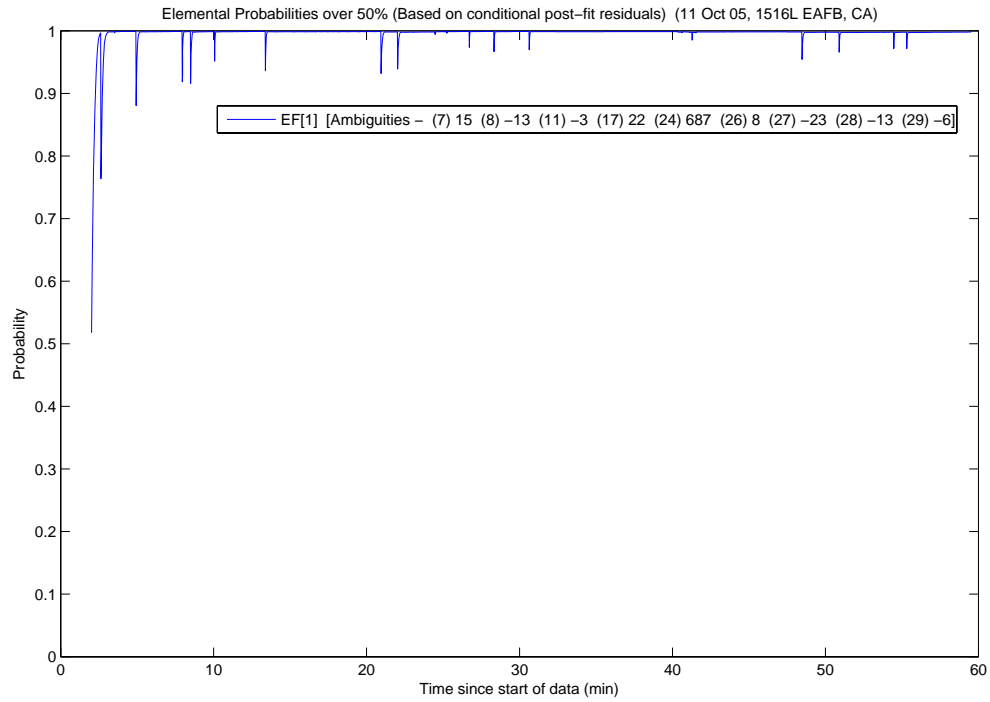


Figure A.213: Case F5.1: Conditional Probabilities for Selected MMAE Elemental Filters

The following figure shows the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.214 is EF[1].

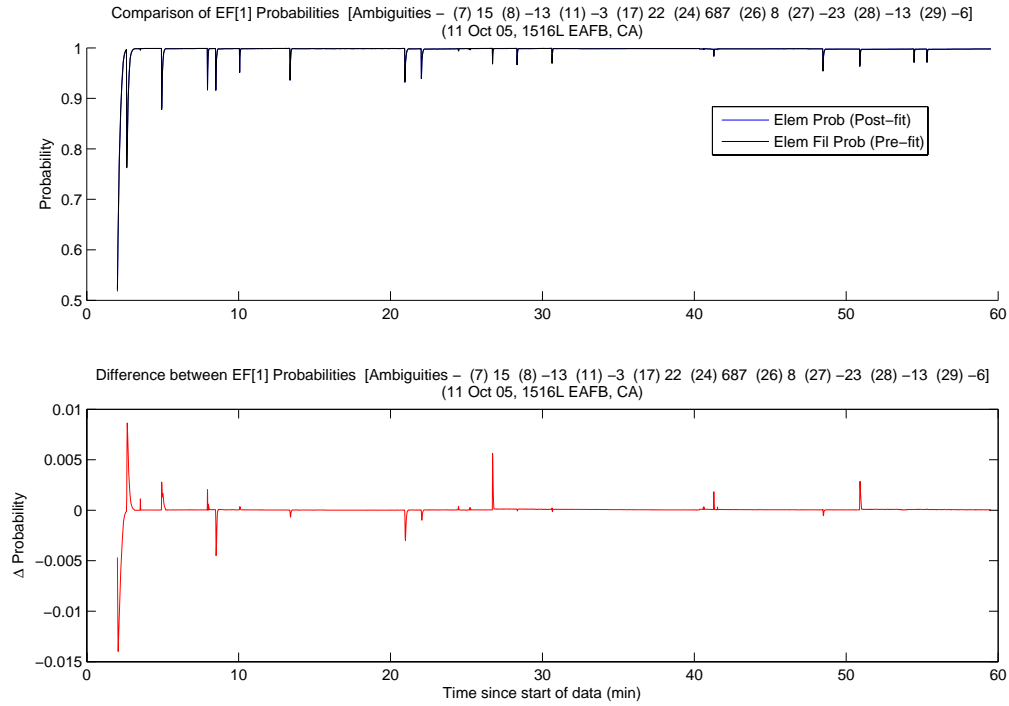


Figure A.214: Case F5.1: EF[1] Probability Comparison

The MMAE position error is shown in Figure A.215. The MMAE solution was almost as good the minimum indicator since the true ambiguity set's probability remained high during the data run. There were a few spikes when an incorrect ambiguity set blended with the true set.

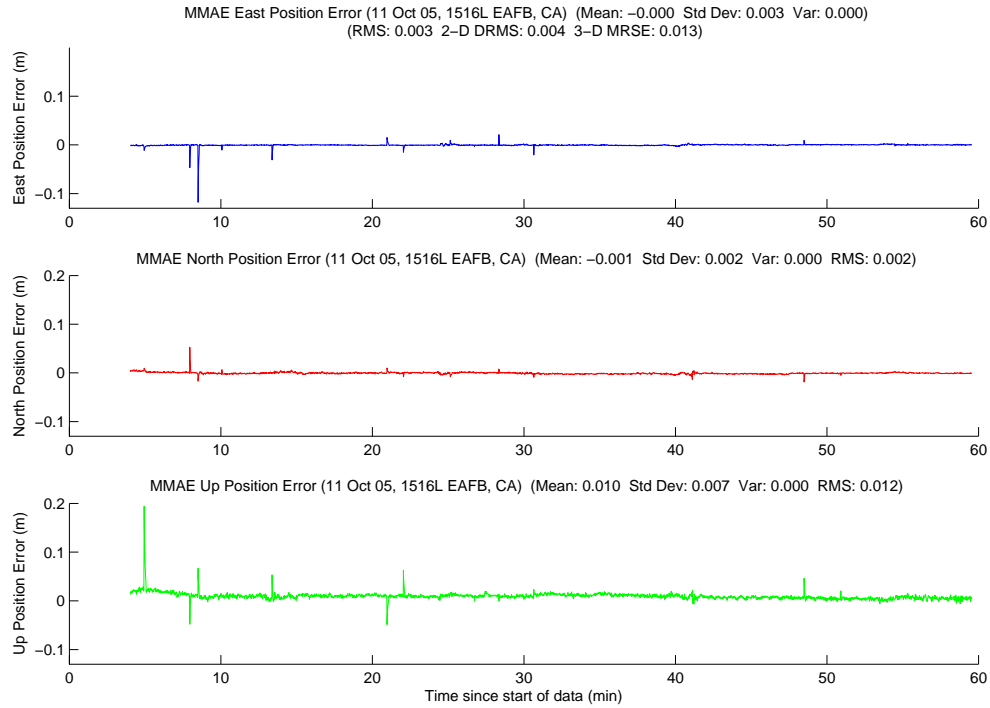


Figure A.215: Case F5.1: MMAE Position Error

Finally, Figure A.216 displays the position error plot for the minimum indicator variable. The MRSE was almost in the millimeter range for the 60 minute data run. The scale on Figure A.216 is four times smaller than the MMAE plot. Lastly, Table A.12 gives a summary of each method's position error.

Table A.12: Case F5.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.045	0.028	-0.036	0.072	0.157	0.064	0.096	0.195
MMAE	-0.000	0.003	-0.001	0.002	0.010	0.007	0.004	0.013
Min. Ind.	-0.000	0.001	-0.001	0.002	0.009	0.005	0.002	0.011

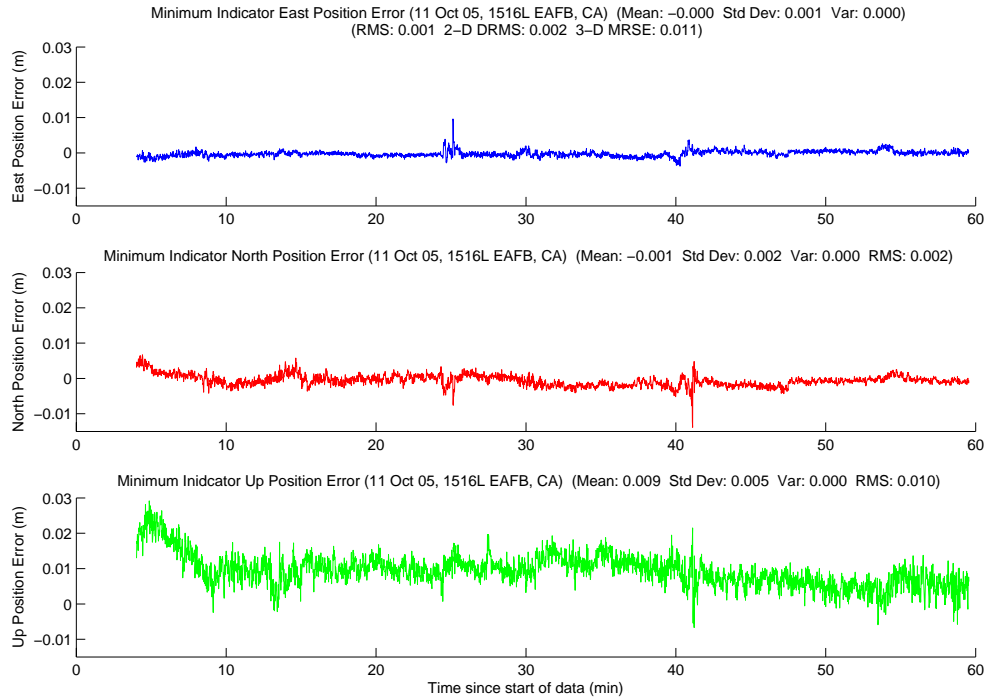


Figure A.216: Case F5.1: Minimum Indicator Position Error

A.2.9 Case F5.2: Flight Test 5, Second Data Set . Case F5.2 started one minute after the end of Case F5.1. The data run was cut short for RTB and lasted 33 minutes, 47 seconds. The aircraft ranged from 10.6 to 4740.4 meters apart during the data run. The 4740.4 meter point occurred at the very end of the data run as the aircraft were separating. Figure A.217 shows the minimum indicator North, East, Down relative positions.

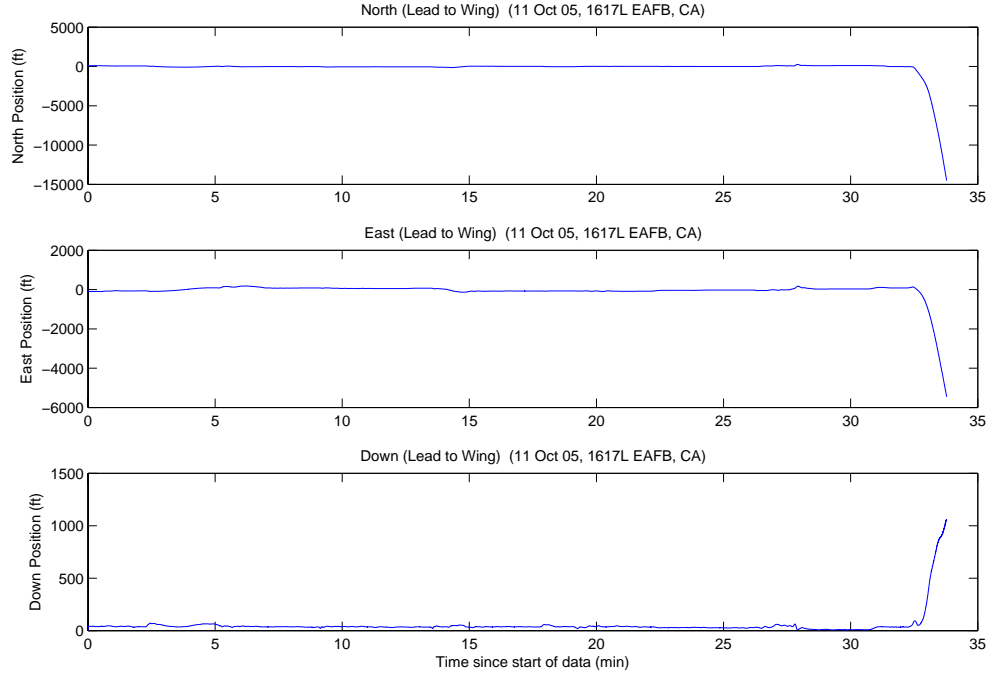


Figure A.217: Case F5.2: North, East, Down Relative Positions

The LAMBDA function returned a total of 31 unique ambiguity sets during the data run. The aircraft were 38.4 meters apart when the system started, with ten visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 1897 of 1907 epochs (99.5 percent).

Figures A.218, A.219, and A.220 show the floating filter errors. The floating filter performed extremely well, as it should with an average of eleven satellites. The floating filter's 1σ estimate of the position error is outside the plot scales in all three figures.

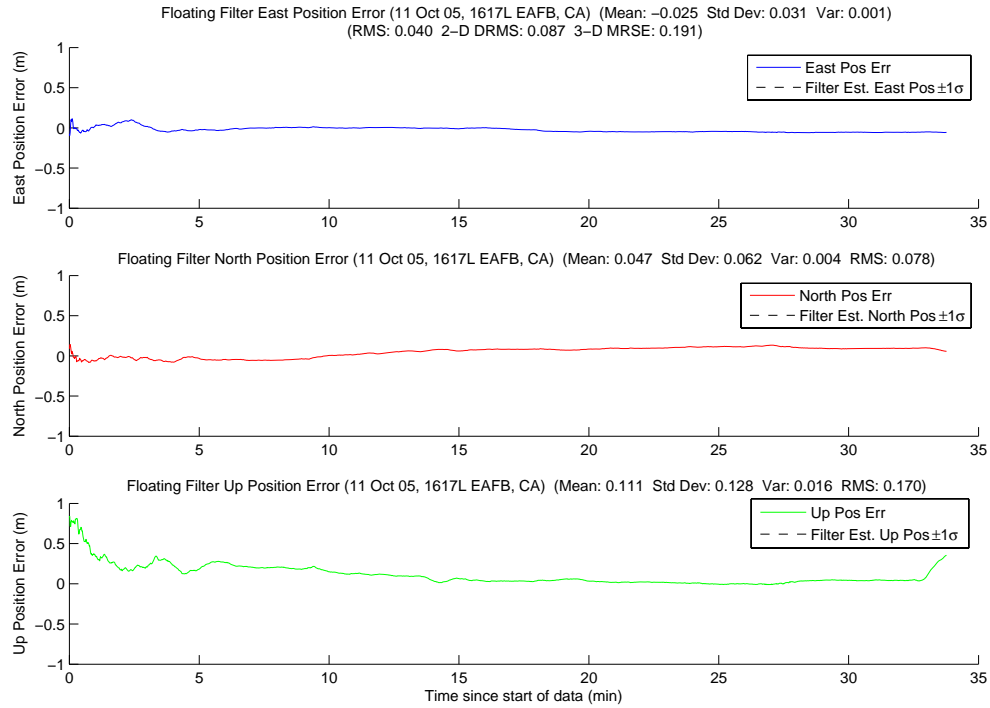


Figure A.218: Case F5.2: Floating Filter Position Errors

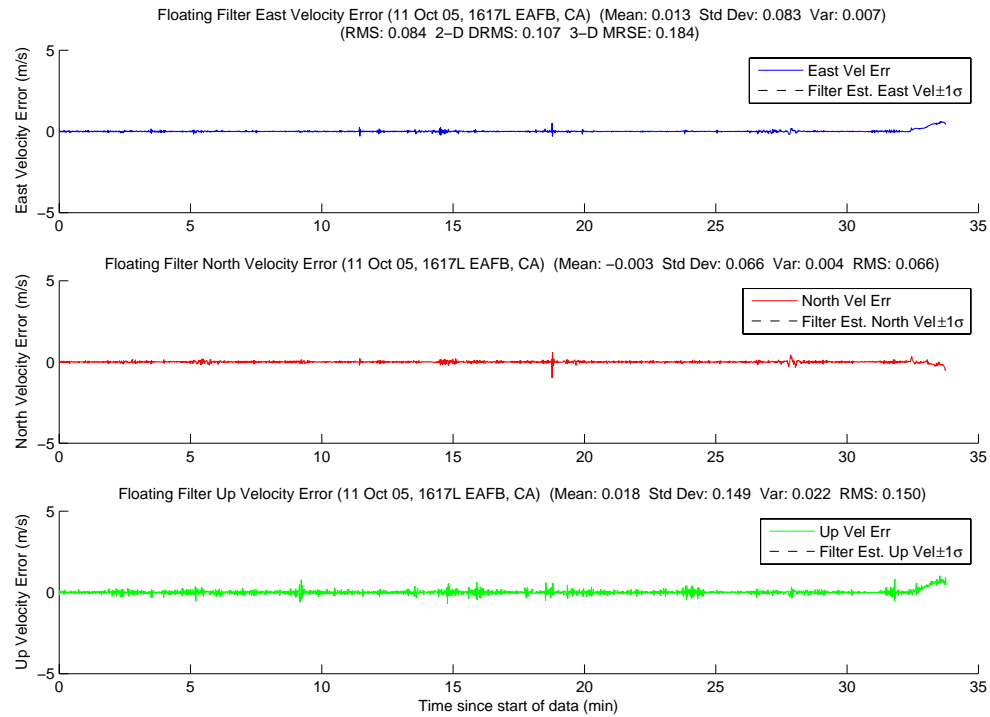


Figure A.219: Case F5.2: Floating Filter Velocity Errors

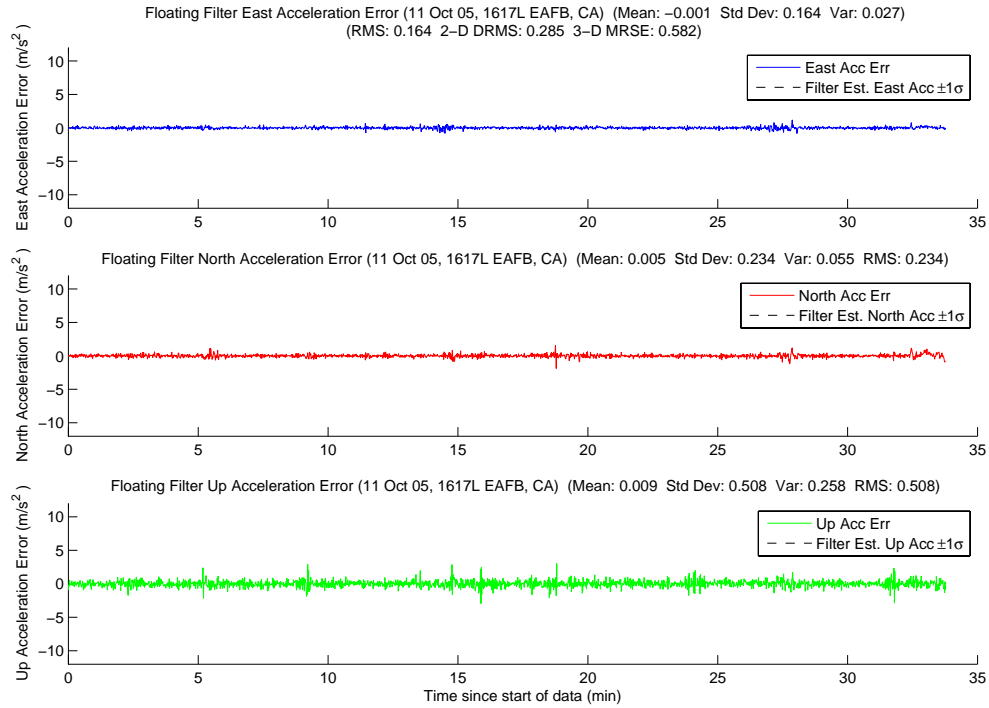


Figure A.220: Case F5.2: Floating Filter Acceleration Errors

Figure A.221 displays PRN 8 measurement information and the floating filter ambiguity estimate. The break in the floating filter 1σ ambiguity estimate at approximately 28 minutes was due to a cycle slip. The rest of the satellite's plots follow in Figures A.222, A.223, A.224, A.225, A.205, A.227, A.228, A.229, and A.230.

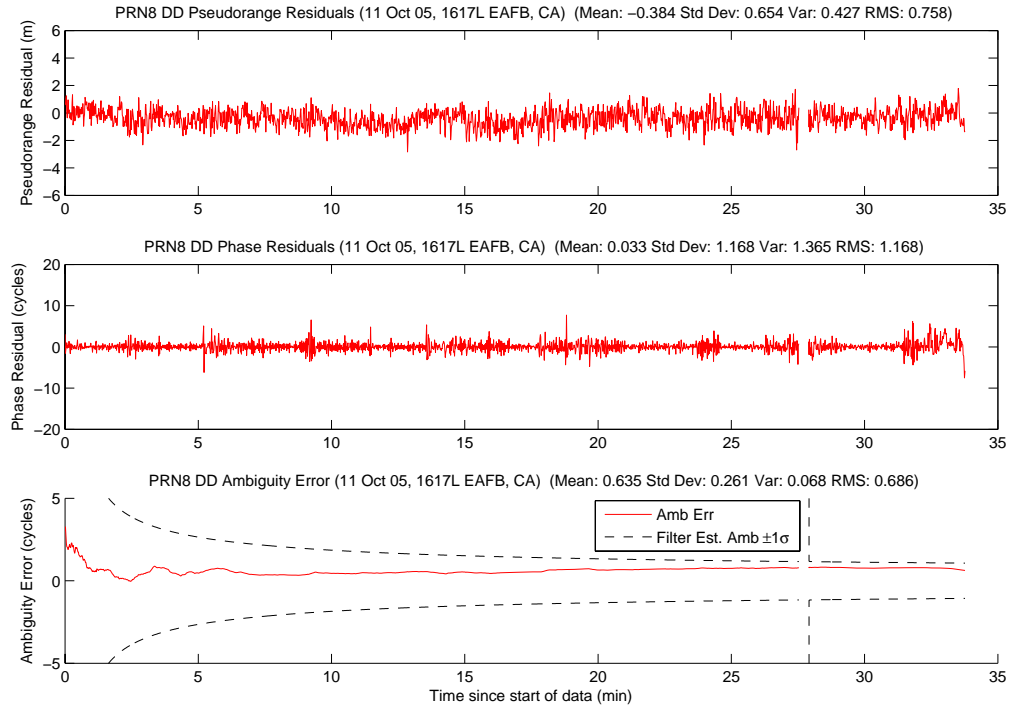


Figure A.221: Case F5.2: Satellite 8 Measurements

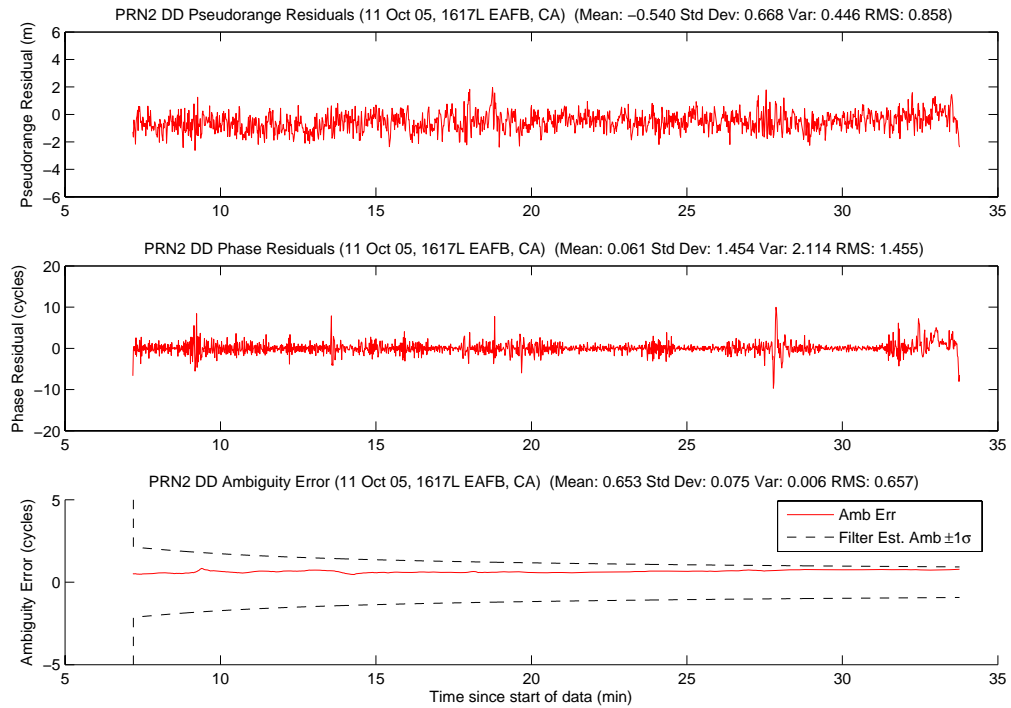


Figure A.222: Case F5.2: Satellite 2 Measurements

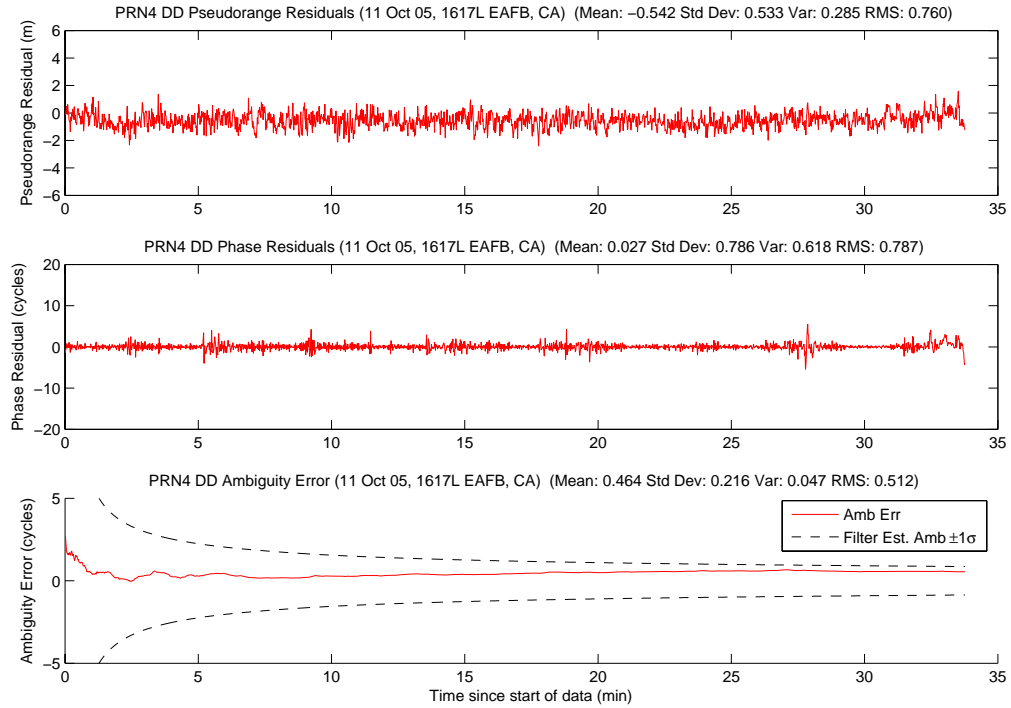


Figure A.223: Case F5.2: Satellite 4 Measurements

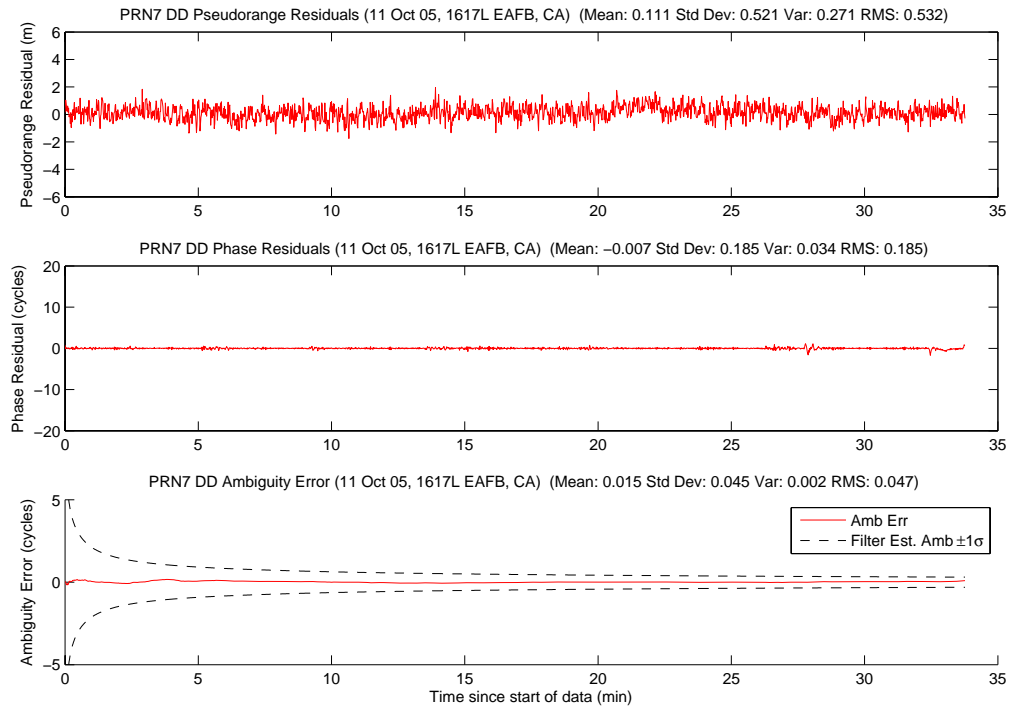


Figure A.224: Case F5.2: Satellite 7 Measurements

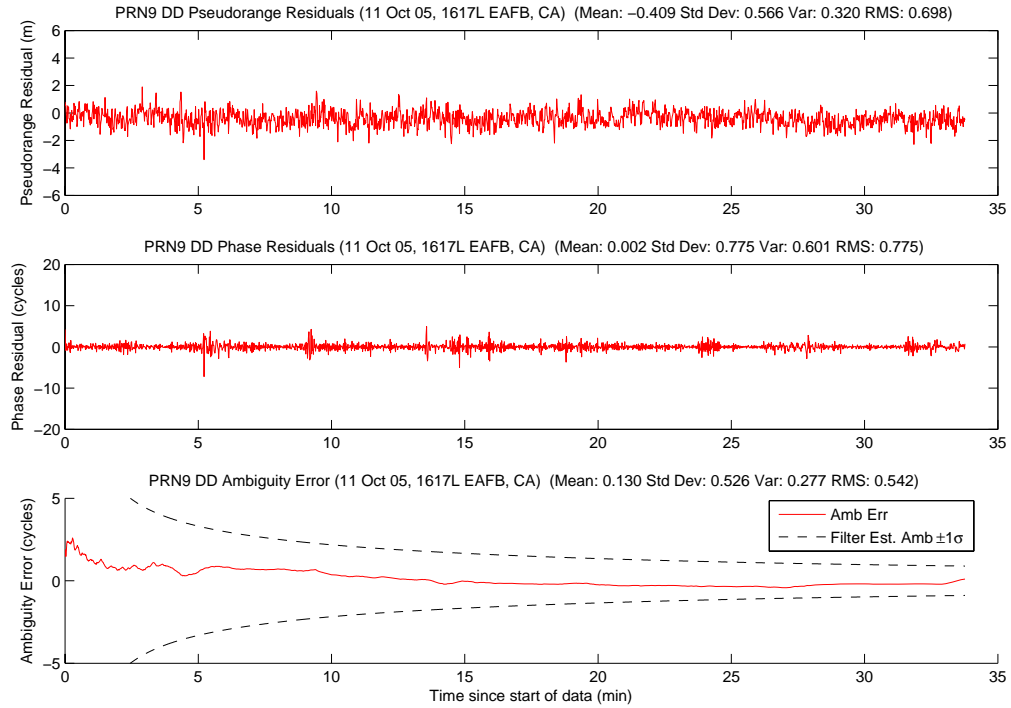


Figure A.225: Case F5.2: Satellite 9 Measurements

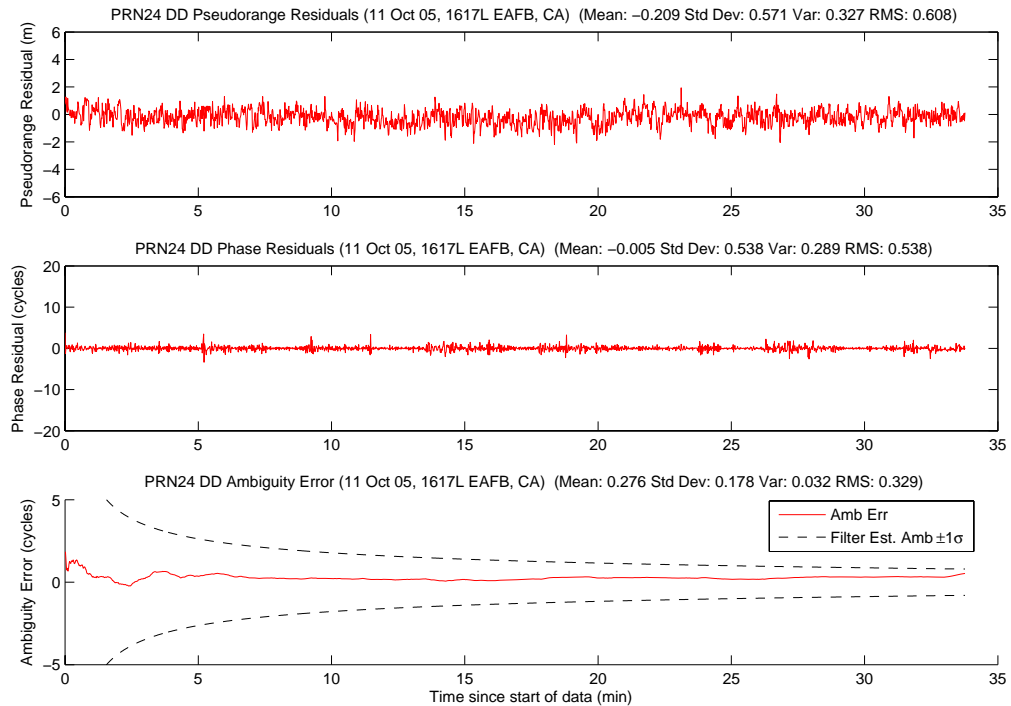


Figure A.226: Case F5.2: Satellite 24 Measurements

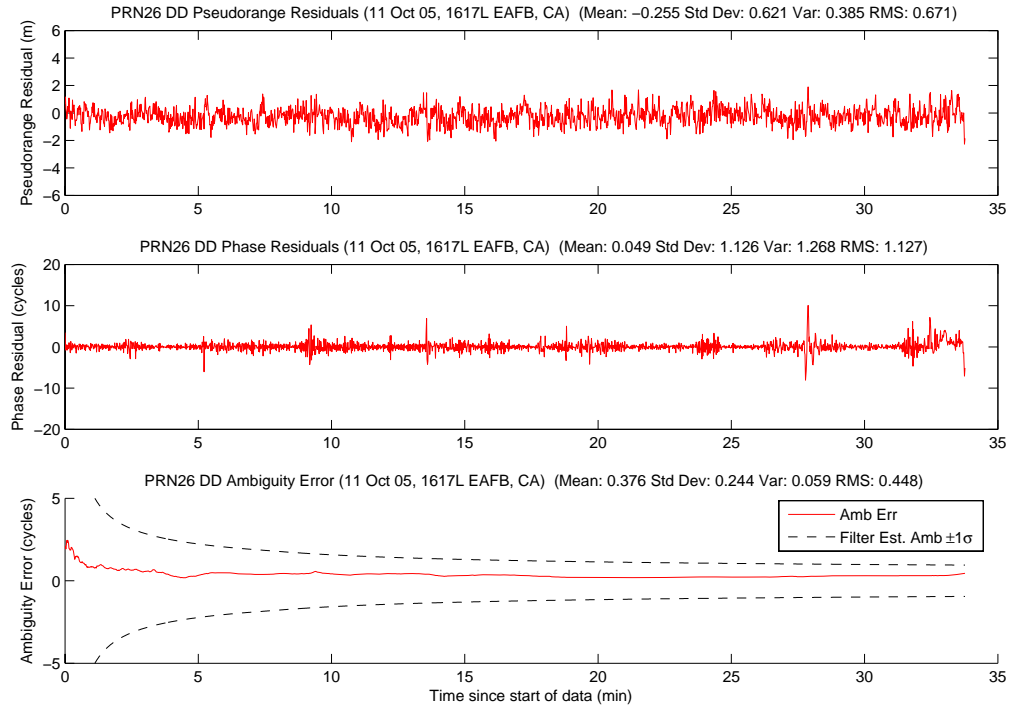


Figure A.227: Case F5.2: Satellite 26 Measurements

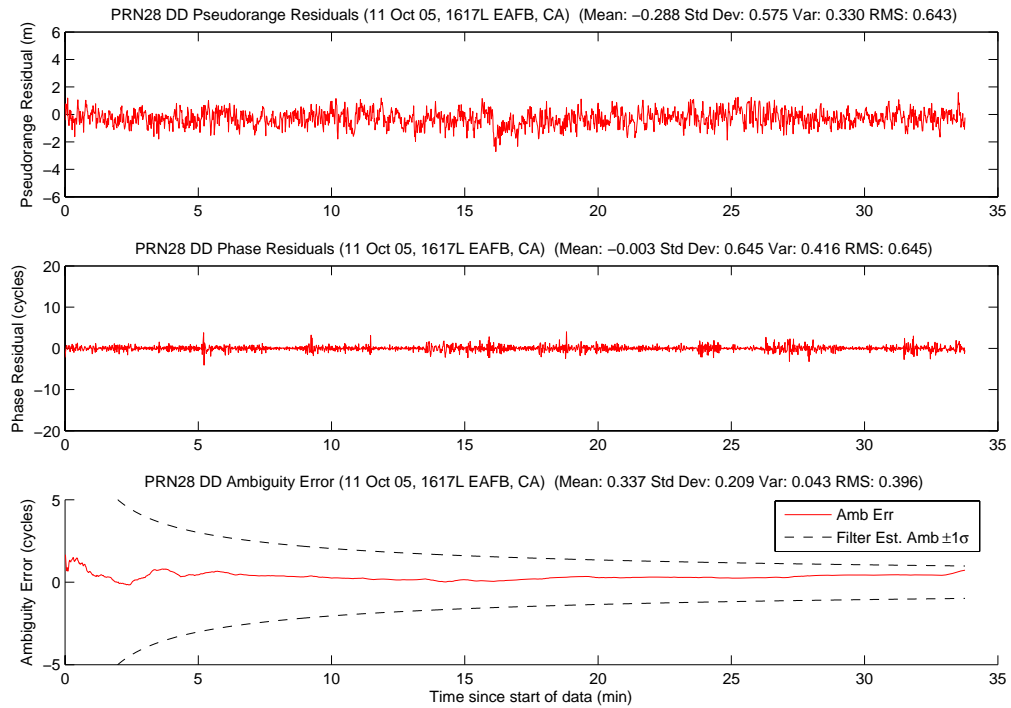


Figure A.228: Case F5.2: Satellite 28 Measurements

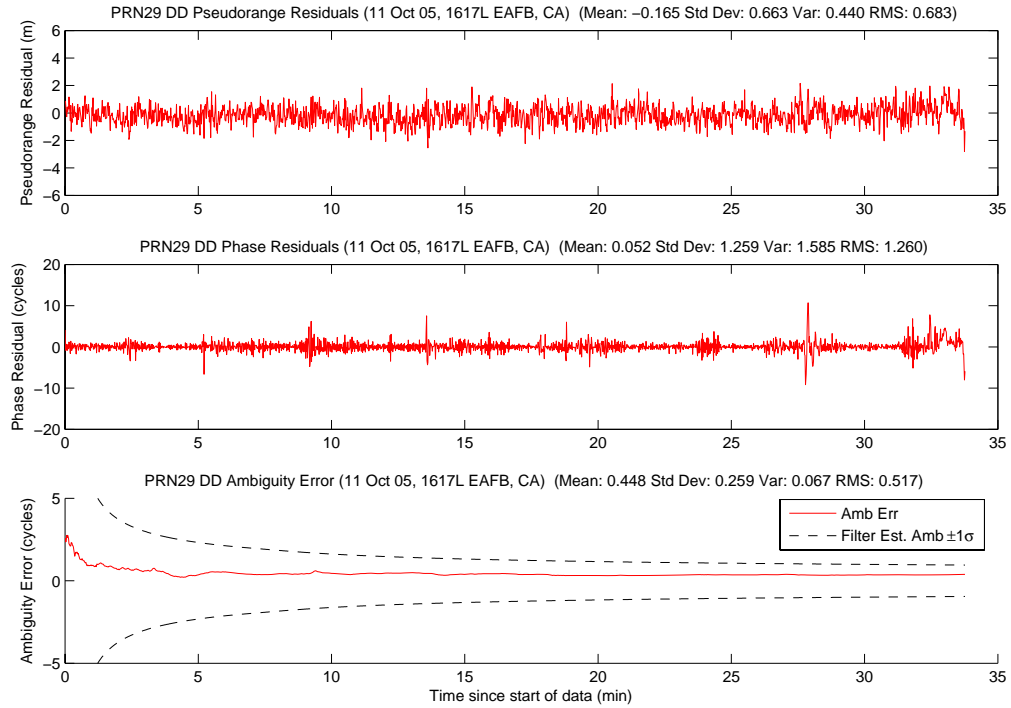


Figure A.229: Case F5.2: Satellite 29 Measurements

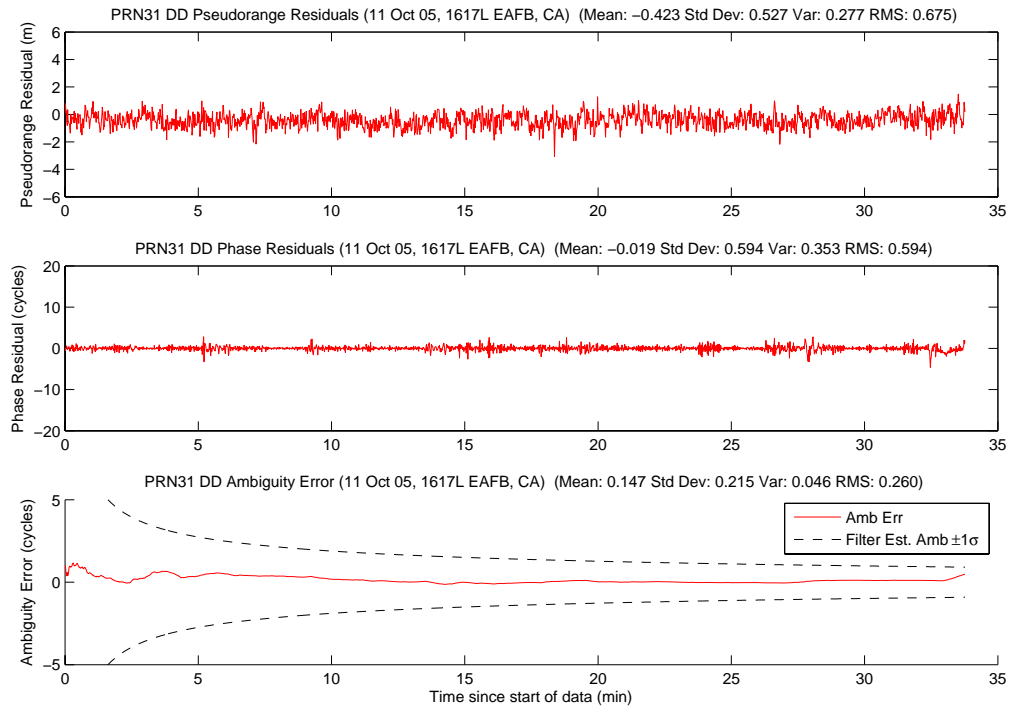


Figure A.230: Case F5.2: Satellite 31 Measurements

Figure A.231 shows the number of visible satellites. The satellite dropout at 28 minutes was due to a cycle slip on PRN 8, as discussed previously.

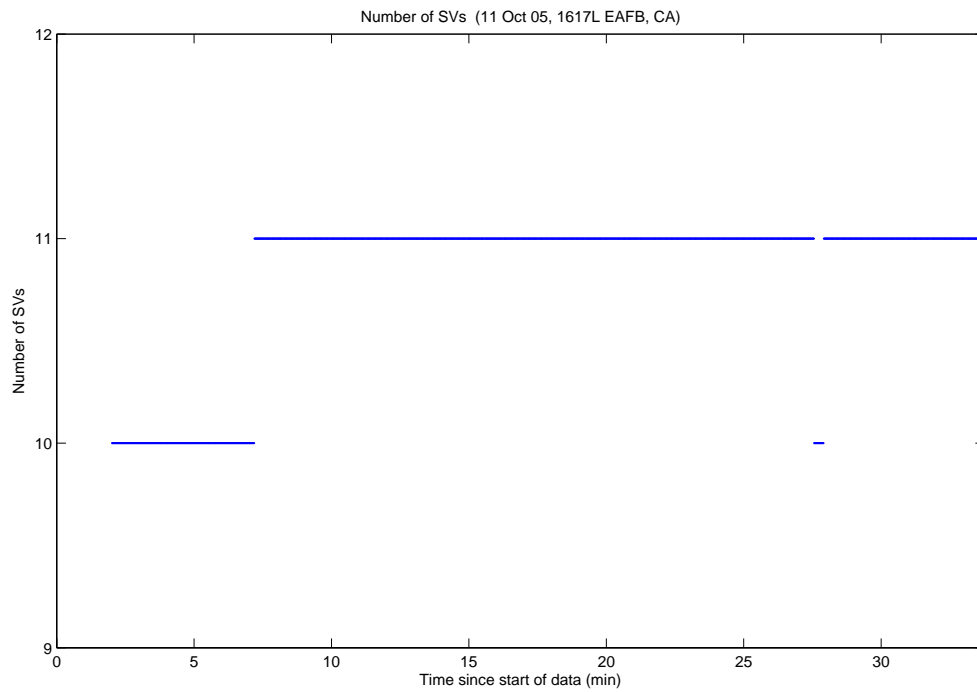


Figure A.231: Case F5.2: Visible Satellites

Figure A.232 shows all of the candidate ambiguity set's SOS residuals. The residual plot is very similar to Case 5.1. There were eleven visible satellites from eight minutes until the end of the data run. Figure A.233 displays the MMAE conditional probability. The true ambiguity set maintained over 49 percent of the total probability for the entire data run after an initial dip at the very beginning of the run. Once there were eleven visible satellites, the true set's probability did not drop below 90 percent.

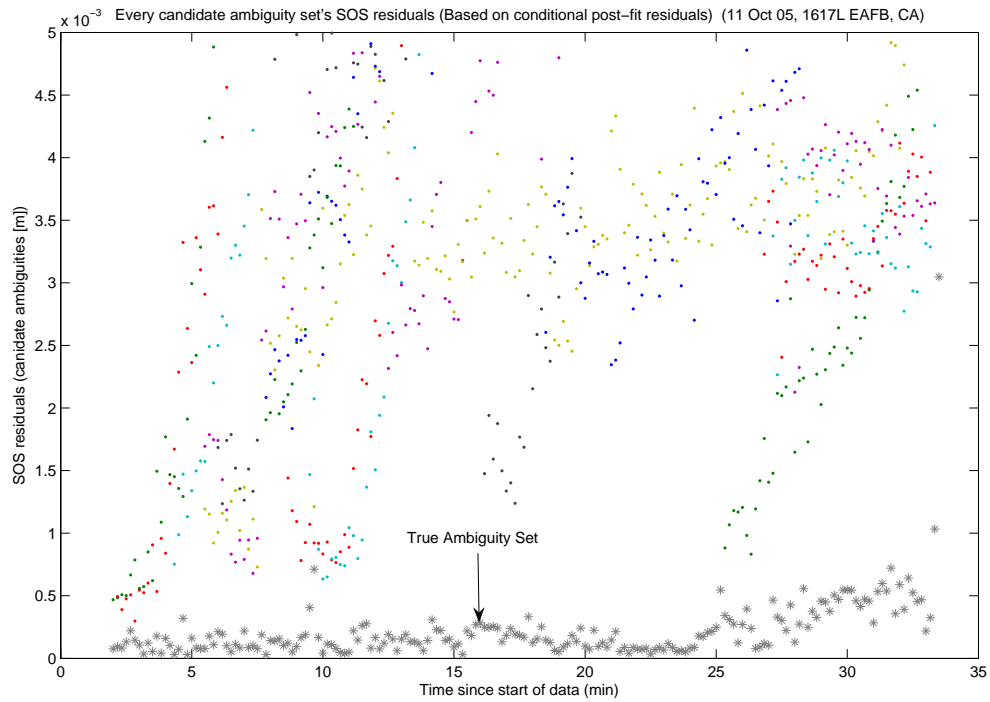


Figure A.232: Case F5.2: SOS Residuals

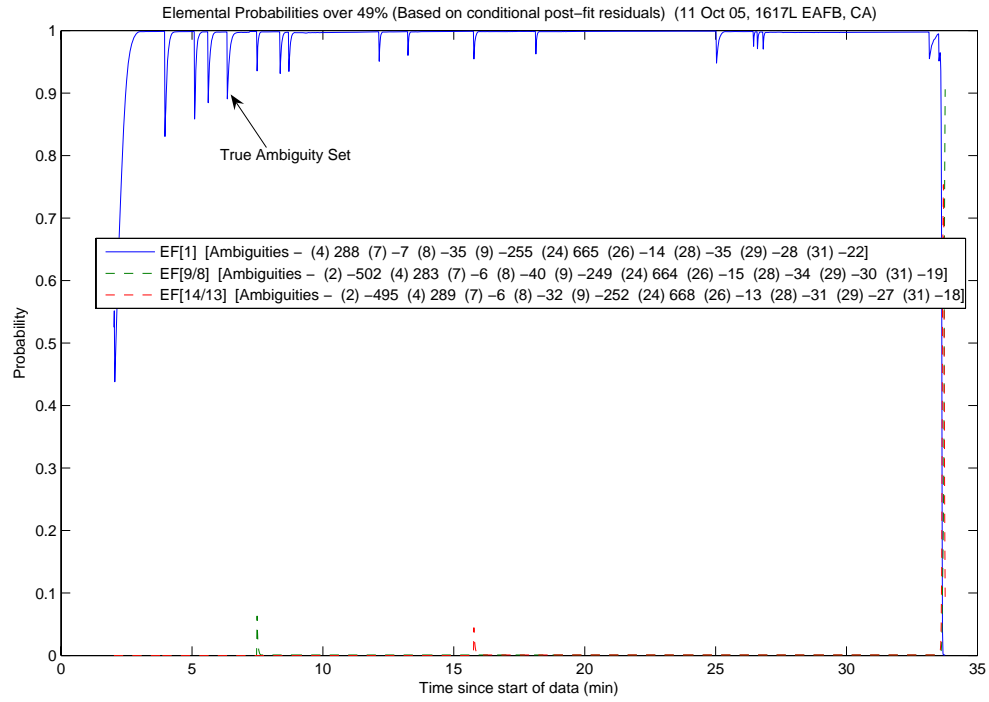


Figure A.233: Case F5.2: Conditional Probabilities for Selected MMAE Elemental Filters

The following figure shows the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.234 is EF[1].

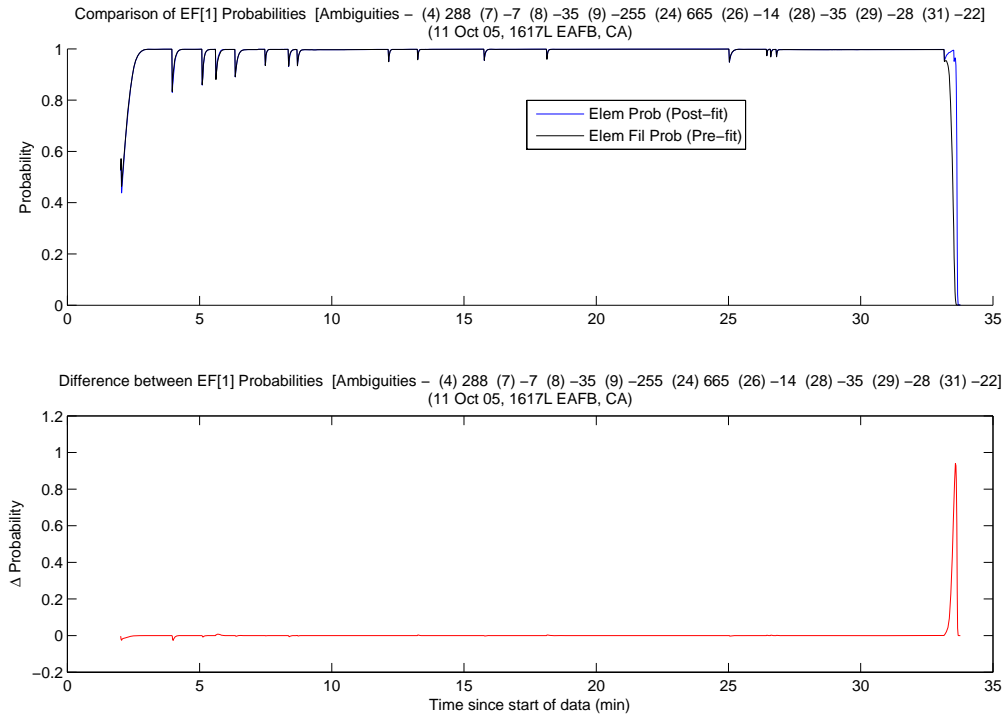


Figure A.234: Case F5.2: EF[1] Probability Comparison

The MMAE position error is shown in Figure A.235. The MMAE solution was almost as good the minimum indicator since the true ambiguity set's probability remained high during the data run. There were a few spikes when an incorrect ambiguity set blended with the true set. Also, at the end of the data run, the MMAE position solution grew considerably due to erroneous ambiguity sets absorbing most of the probability.

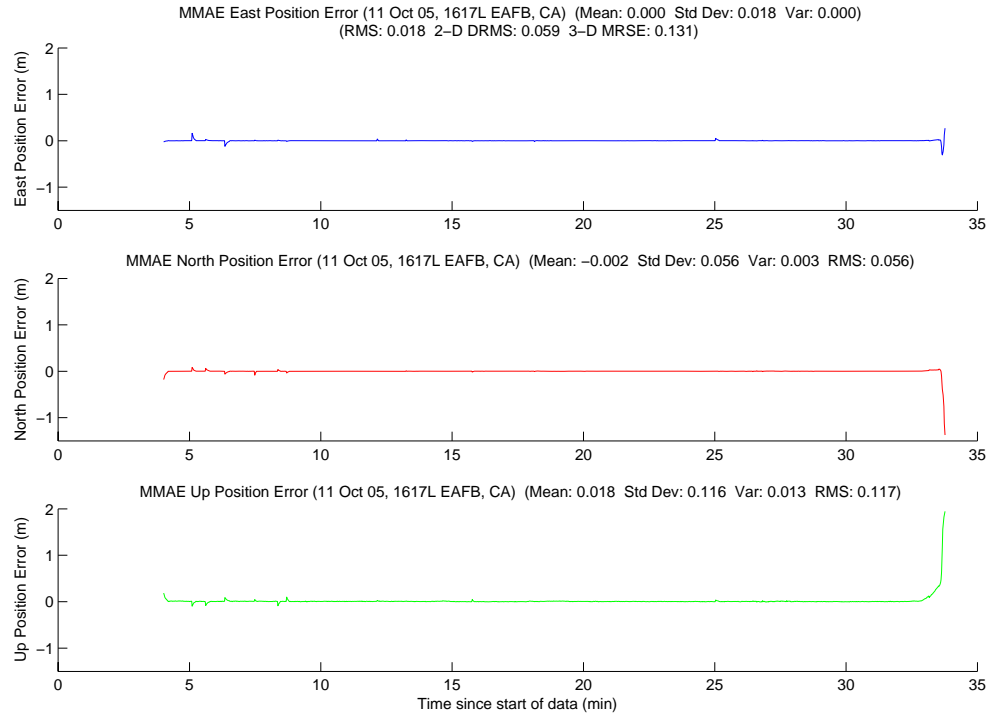


Figure A.235: Case F5.2: MMAE Position Error

Finally, Figure A.236 displays the position error plot for the minimum indicator variable. Since the minimum indicator was “locked” to the true ambiguity set for the entire run, the position errors were very small. Even when the aircraft were almost 5 kilometers apart, the North and East errors remain extremely low. The largest contributor to the total error came from the Up error. In fact, the DRMS was 6 millimeters for the entire data run! Lastly, Table A.13 gives a summary of each method’s position error.

Table A.13: Case F5.2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.025	0.031	0.047	0.062	0.111	0.128	0.087	0.191
MMAE	0.000	0.018	-0.002	0.056	0.018	0.116	0.059	0.131
Min. Ind.	0.001	0.004	0.002	0.004	0.013	0.049	0.006	0.050

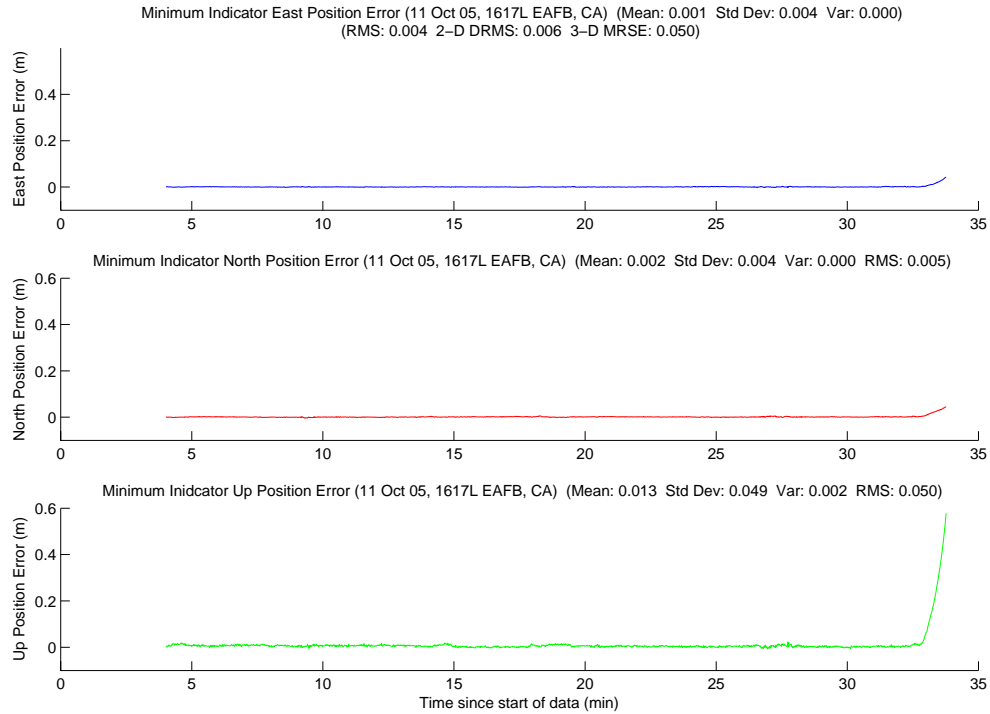


Figure A.236: Case F5.2: Minimum Indicator Position Error

A.2.10 Case F6.1: First Flight of Flight Test 6, First Data Set . There were two test flights on 12 October 2005. Each flight had two data runs for a total of four runs. The first data run started at 0918L and lasted 59 minutes, 33 seconds. The aircraft ranged from 12.3 to 156.1 meters apart during the data run. Figure A.237 shows the minimum indicator North, East, Down relative positions.

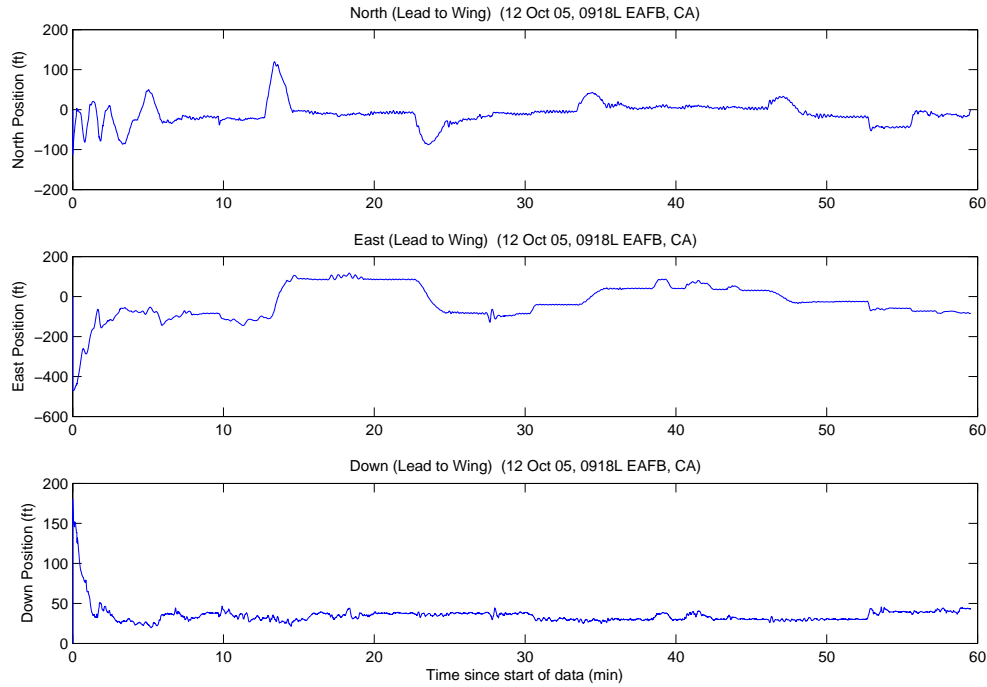


Figure A.237: Case F6.1: North, East, Down Relative Positions

A total of 30 unique ambiguity sets were returned from the LAMBDA function during the data run. The aircraft were 156 meters apart when the system started, with eight visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 3453 of 3453 epochs (100.0 percent), with a maximum ratio of 74.7.

Figures A.238, A.239, and A.240 show the floating filter position errors. The floating filter performed very well. The floating filter's 1σ estimate of the position error is outside the plot scales in all three figures.

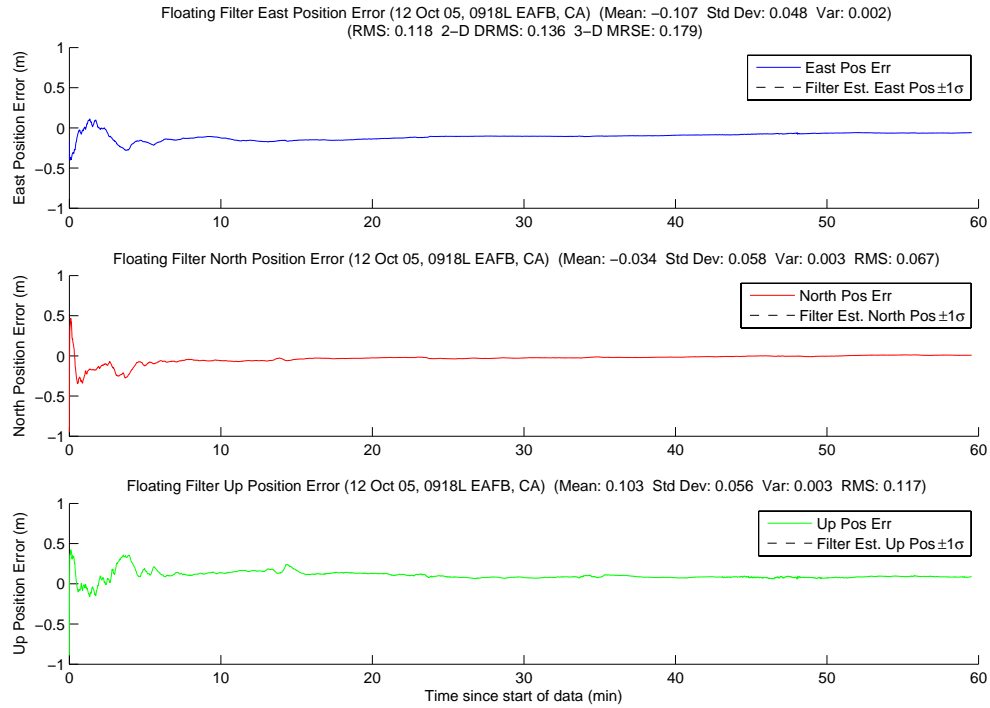


Figure A.238: Case F6.1: Floating Filter Position Errors

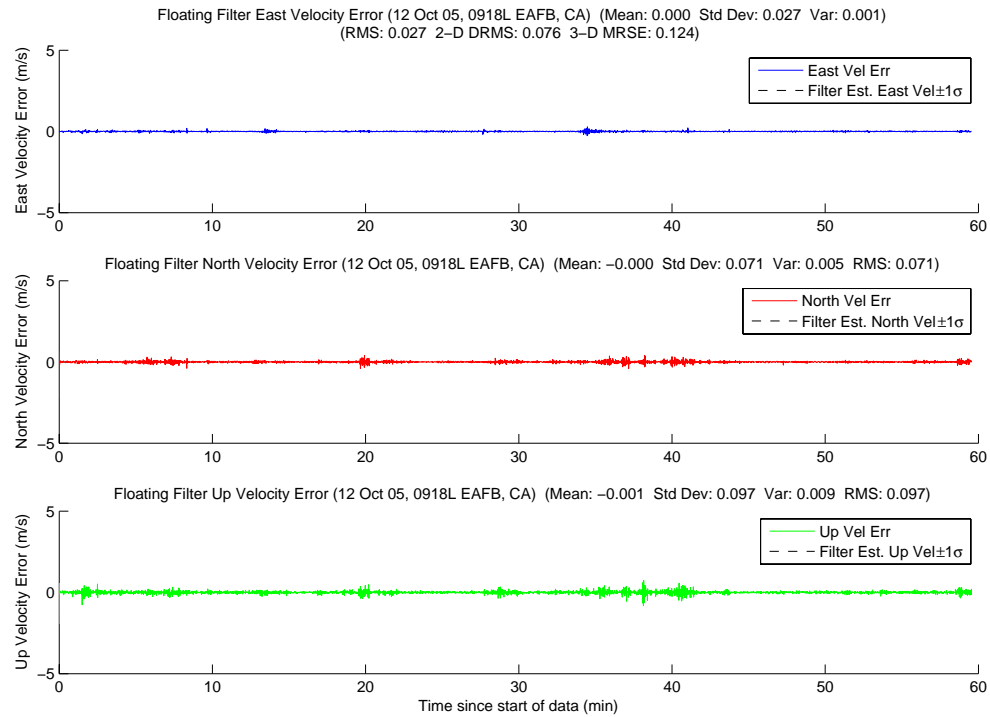


Figure A.239: Case F6.1: Floating Filter Velocity Errors

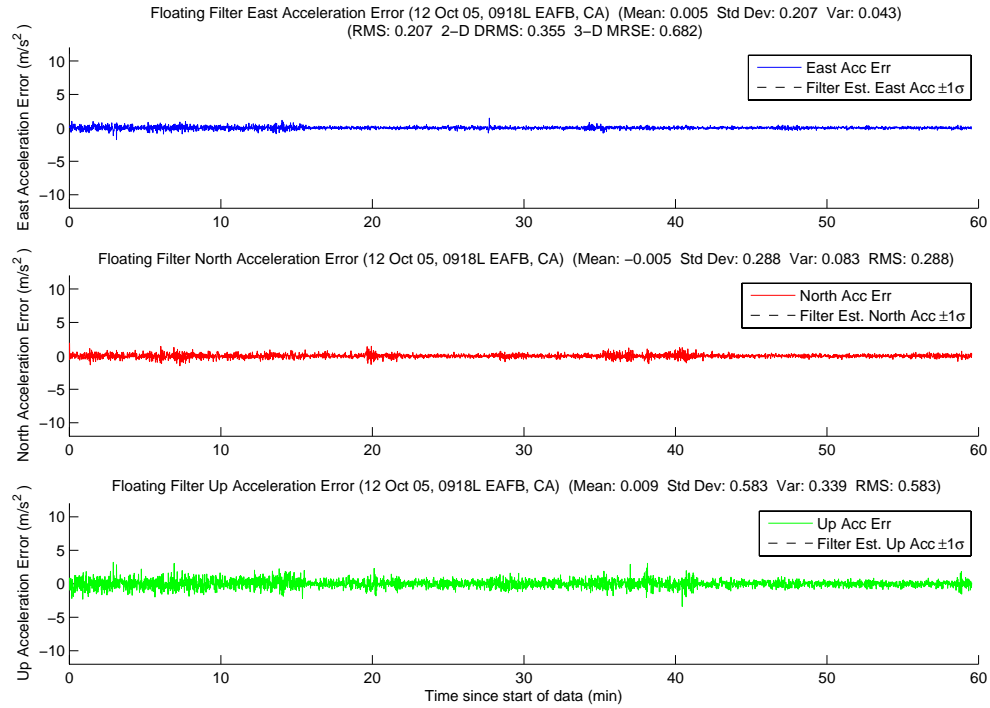


Figure A.240: Case F6.1: Floating Filter Acceleration Errors

Figure A.241 displays PRN 16 measurement information and the floating filter ambiguity estimate. The rest of the satellite's plots follow in Figures A.242, A.243, A.244, A.245, A.246, A.247, A.248, and A.249.

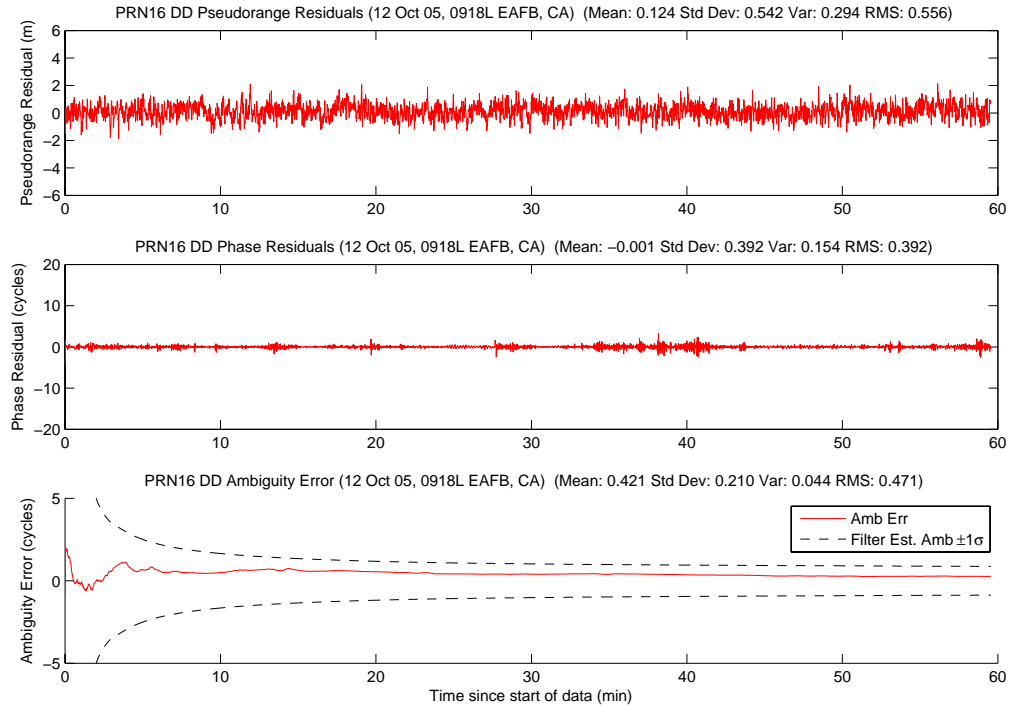


Figure A.241: Case F6.1: Satellite 16 Measurements

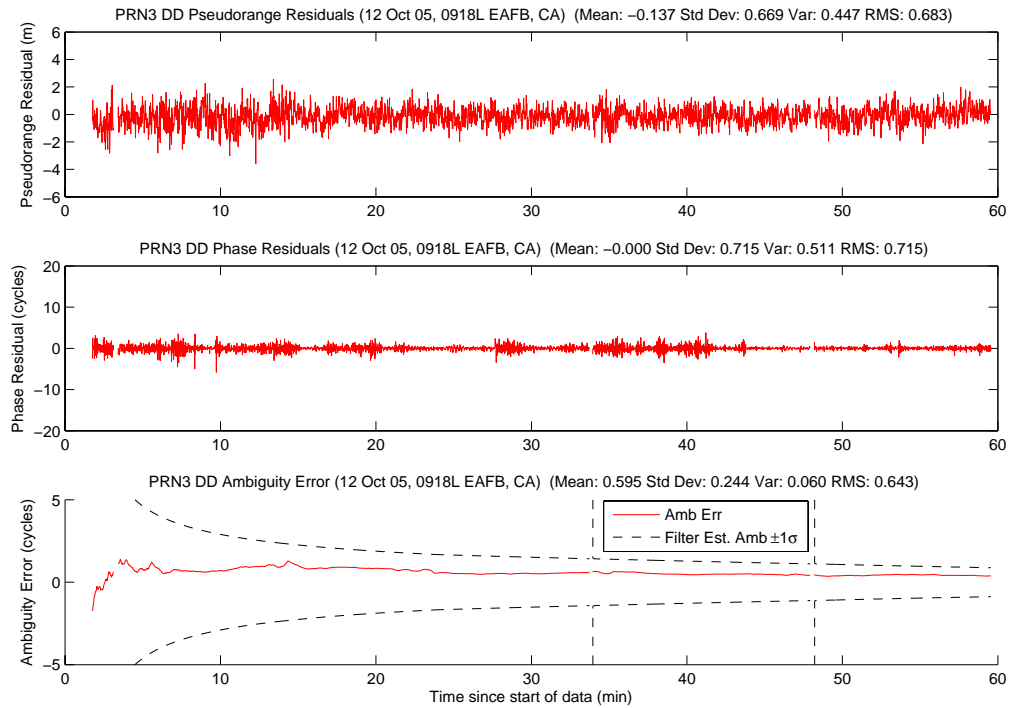


Figure A.242: Case F6.1: Satellite 3 Measurements

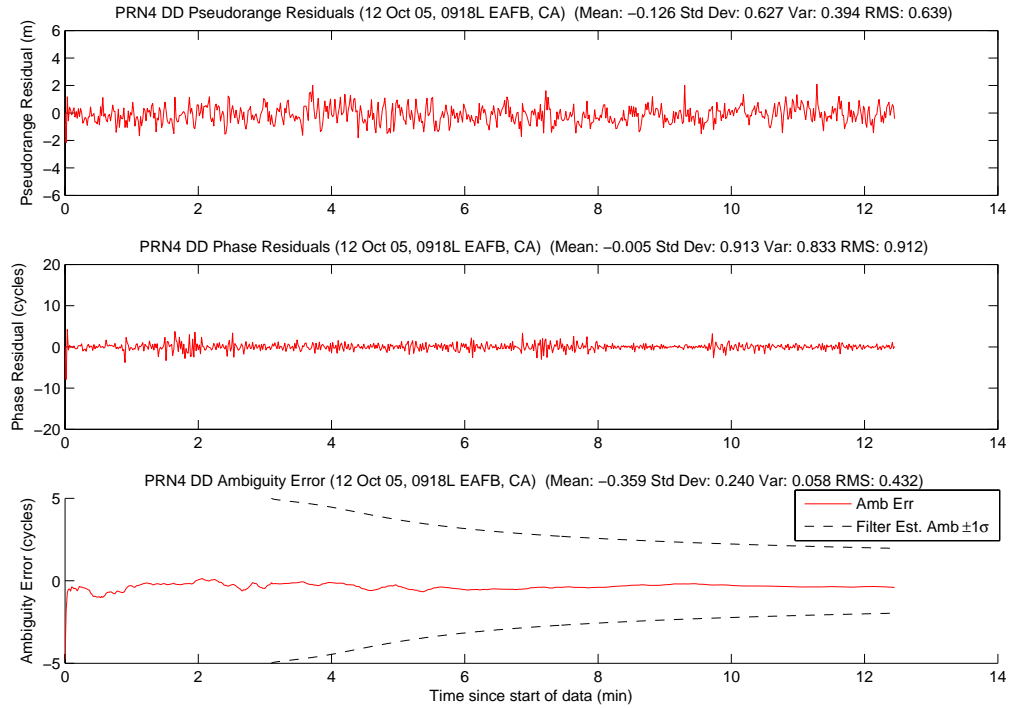


Figure A.243: Case F6.1: Satellite 4 Measurements

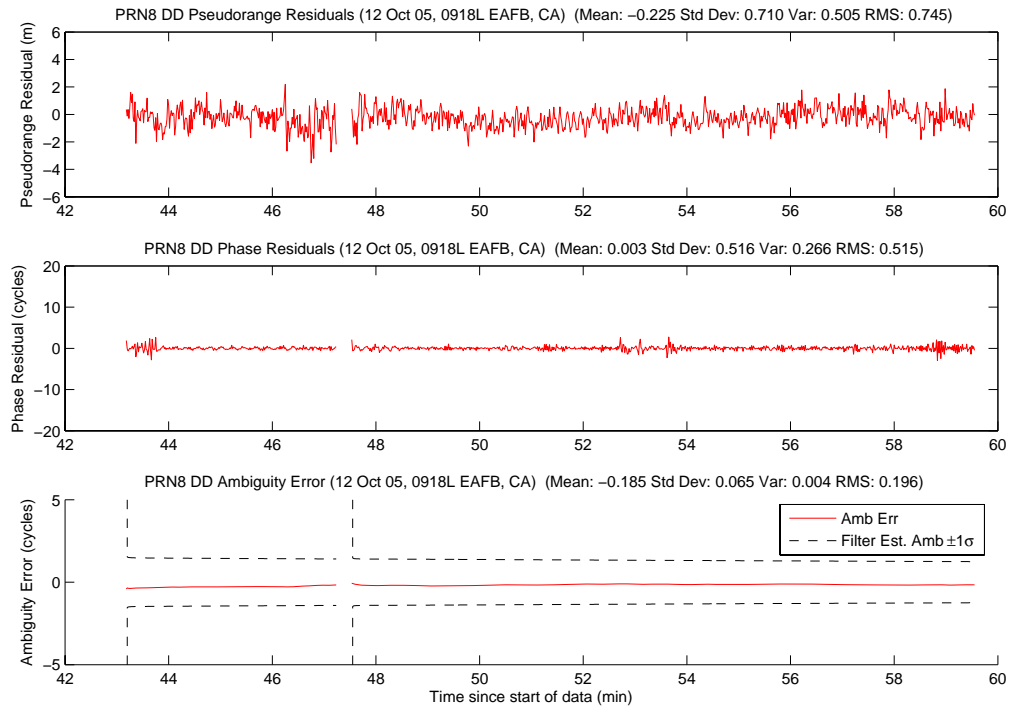


Figure A.244: Case F6.1: Satellite 8 Measurements

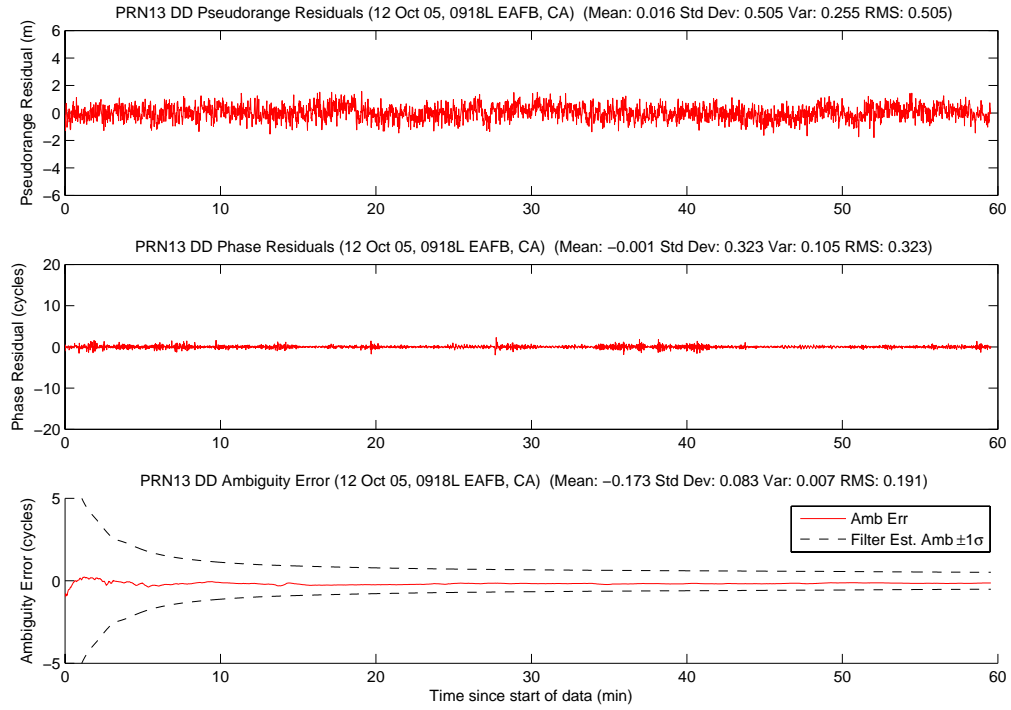


Figure A.245: Case F6.1: Satellite 13 Measurements

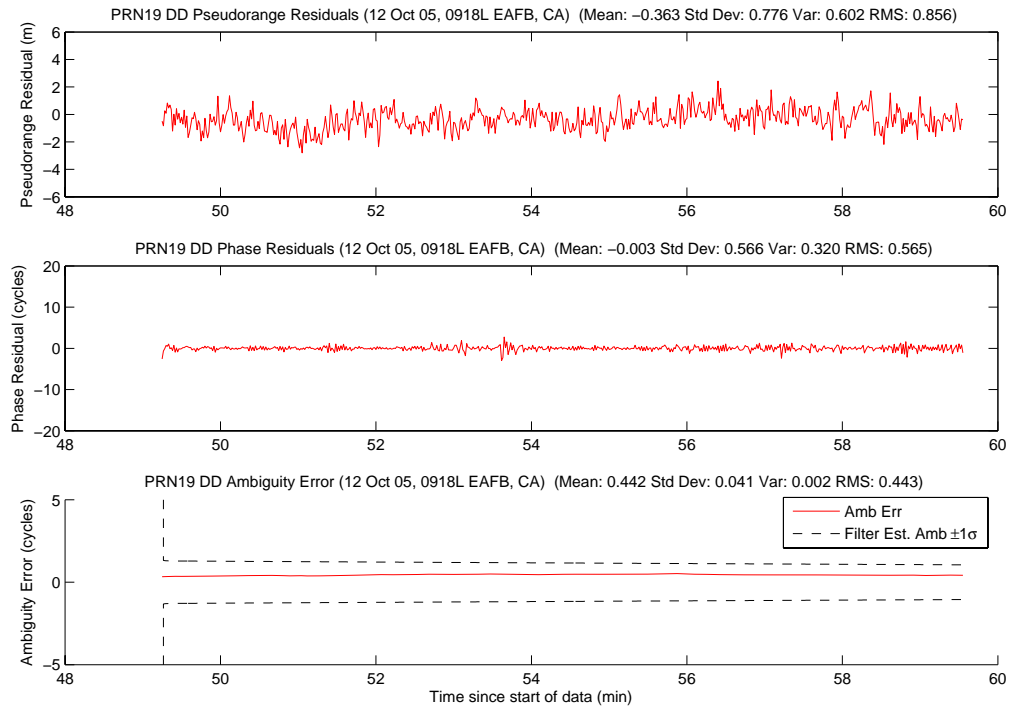


Figure A.246: Case F6.1: Satellite 19 Measurements

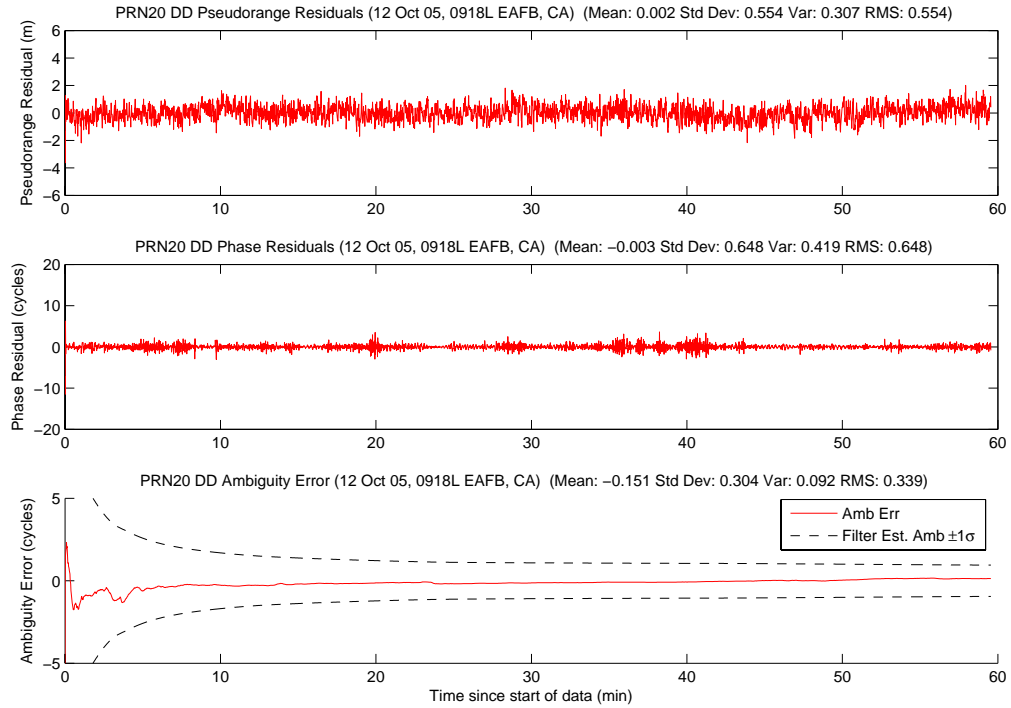


Figure A.247: Case F6.1: Satellite 20 Measurements

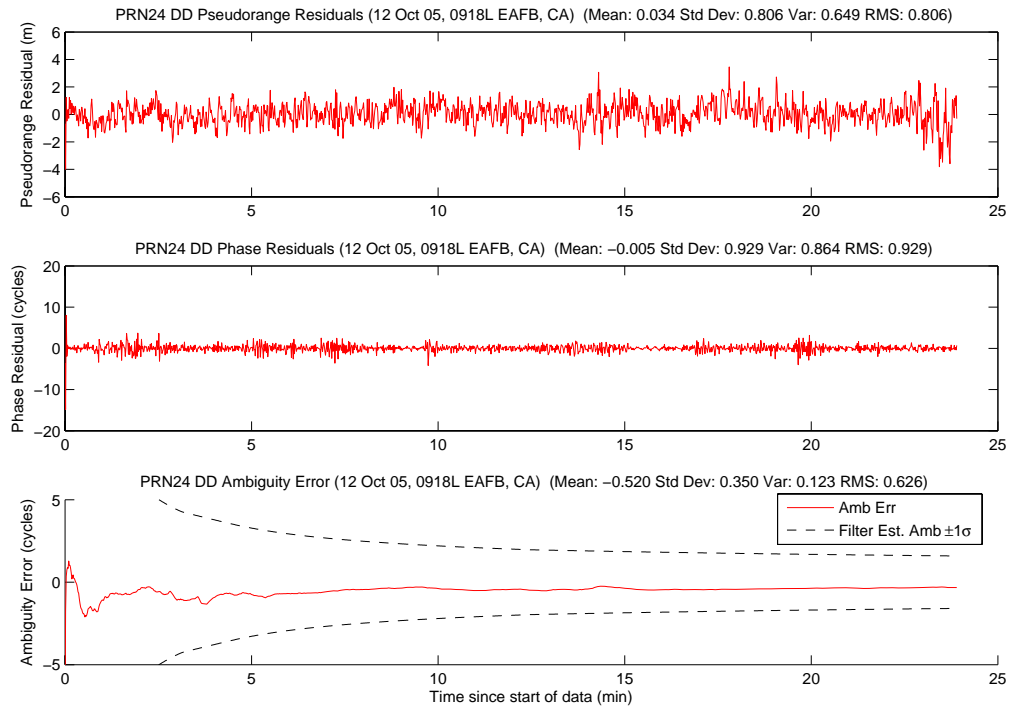


Figure A.248: Case F6.1: Satellite 24 Measurements

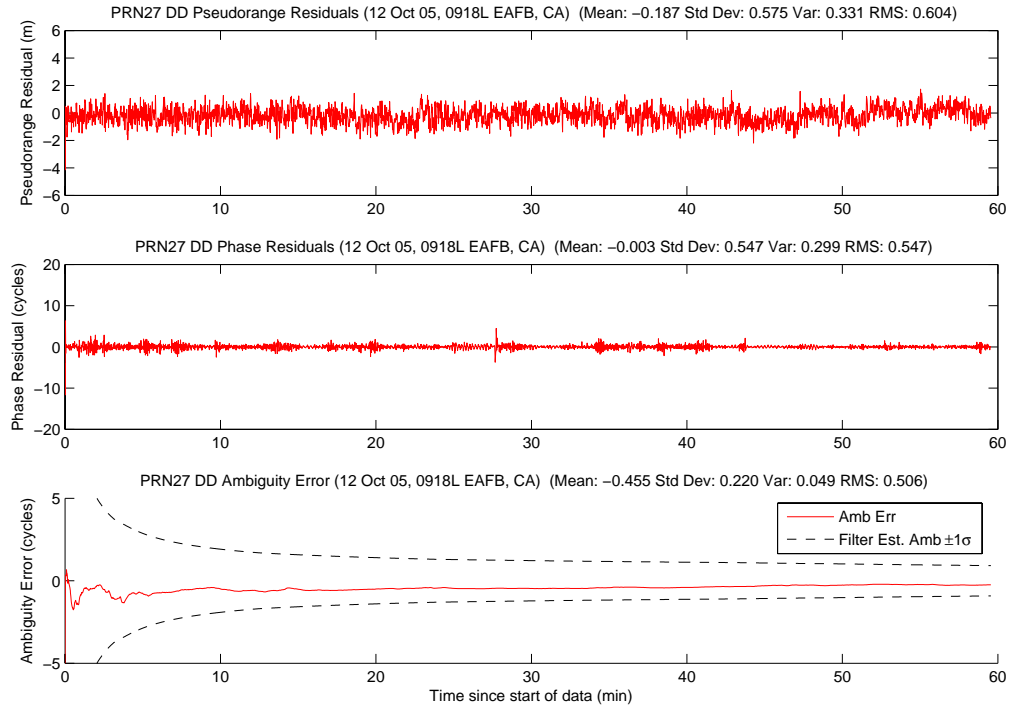


Figure A.249: Case F6.1: Satellite 27 Measurements

Figure A.250 shows the number of visible satellites. All of the satellite dropouts were due to satellite masking. Masking of PRN 3 occurred at 3, 34, and 48 minutes. PRN 8 masked at 47 minutes, four minutes after it had risen. PRN 4 set at 12 minutes and PRN 25 set at 24 minutes. Later in the data run, PRN 8 rose at 43 minutes and PRN 19 rose at 49 minutes.

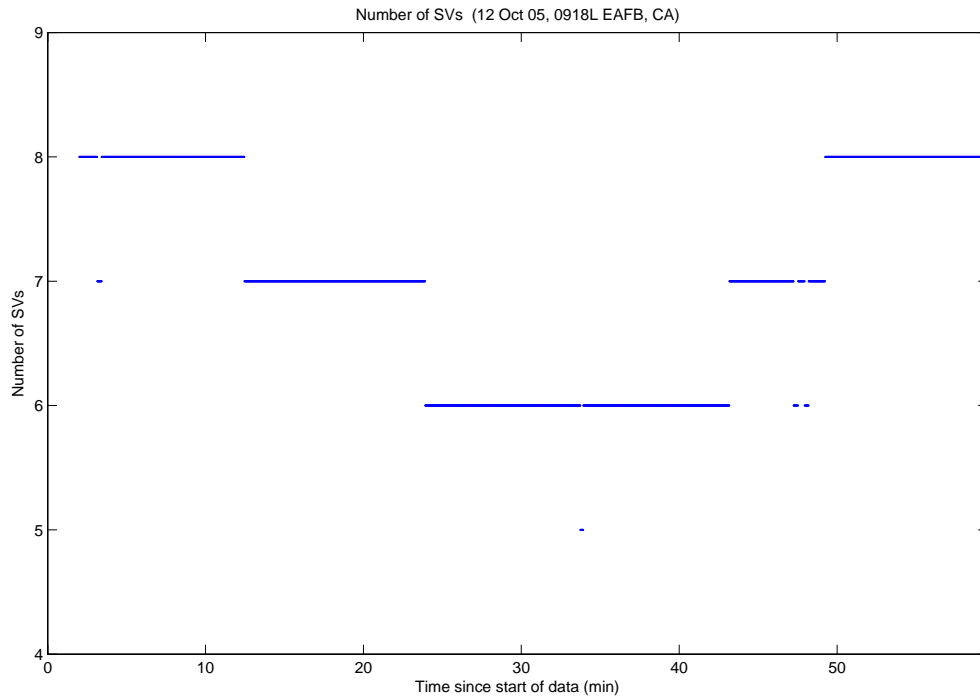


Figure A.250: Case F6.1: Visible Satellites

Figure A.251 shows all of the candidate ambiguity set's SOS residuals. In stark contrast to the residual plots from Cases 5.1 and 5.2, there were many erroneous ambiguity set's residuals near or below the true ambiguity set. This was especially true during the middle of the data run when there were only six visible satellites. Figure A.252 displays the MMAE conditional probability. At 10, 32, and 40 minutes, there were three large dips in the true ambiguity set's probability. At each location in the SOS residuals plot, there was an erroneous ambiguity set's residuals clearly below the true set's residuals. At 10 minutes, the erroneous set's residuals are open green squares; at 32 minutes, the residuals are open red diamonds; at 40 minutes, the residuals are open gold triangles.

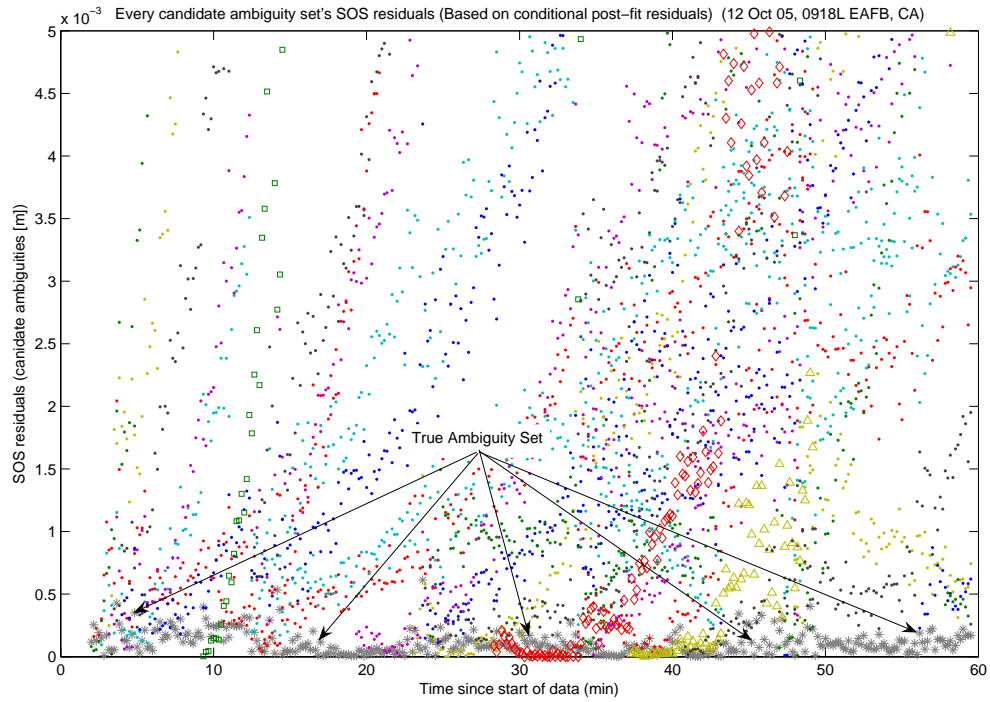


Figure A.251: Case F6.1: SOS Residuals

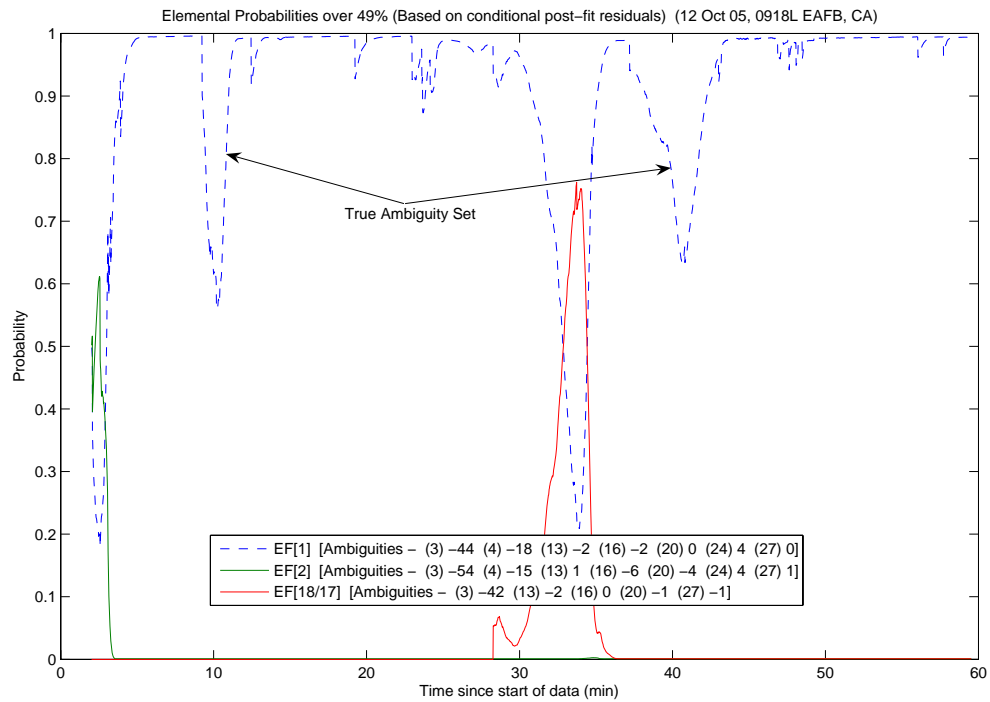


Figure A.252: Case F6.1: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.253 is EF[1], Figure A.254 is EF[2], and Figure A.255 is EF[18].

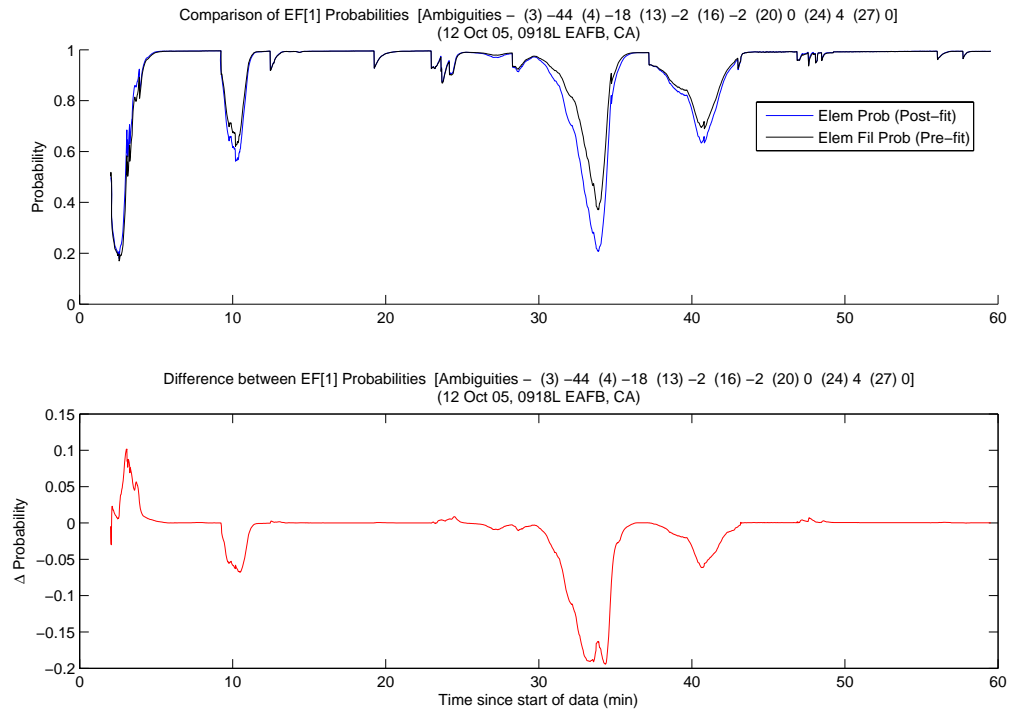


Figure A.253: Case F6.1: EF[1] Probability Comparison

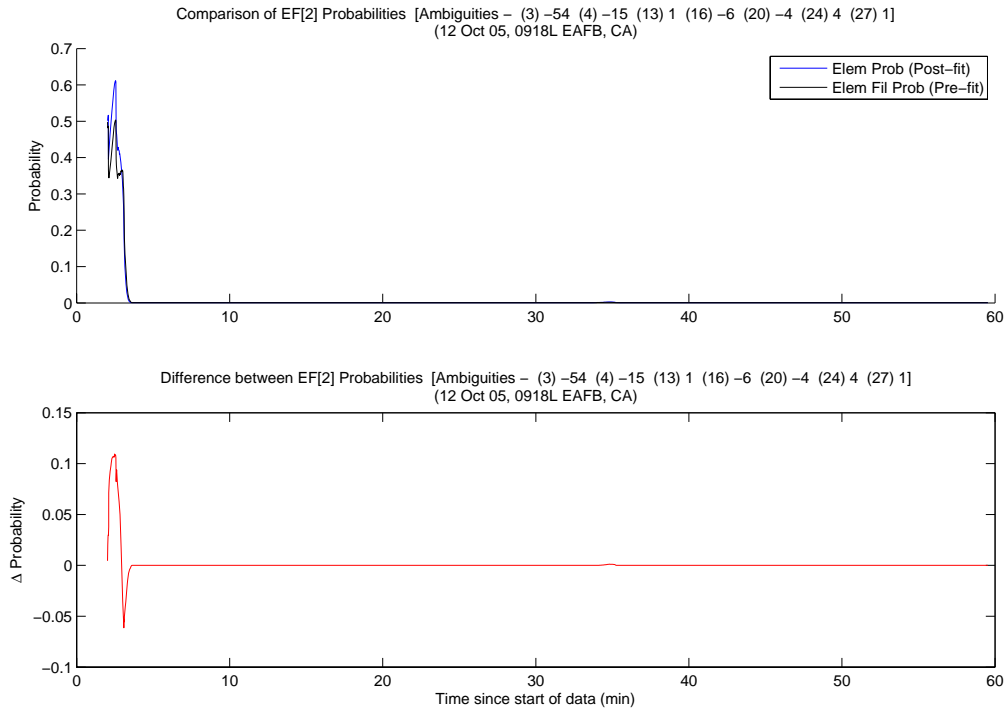


Figure A.254: Case F6.1: EF[2] Probability Comparison

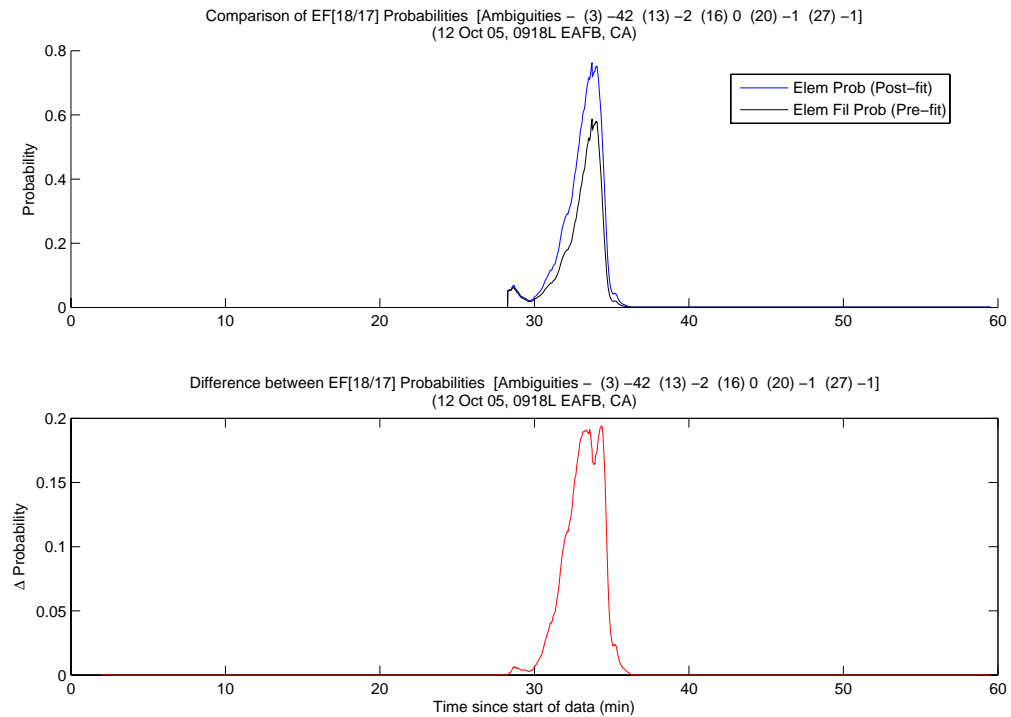


Figure A.255: Case F6.1: EF[18] Probability Comparison

The MMAE position error is shown in Figure A.256. The times described above, where an erroneous set's solution was blended with the true set's solution, are clearly evident. Otherwise, the MMAE position solution was generally accurate.

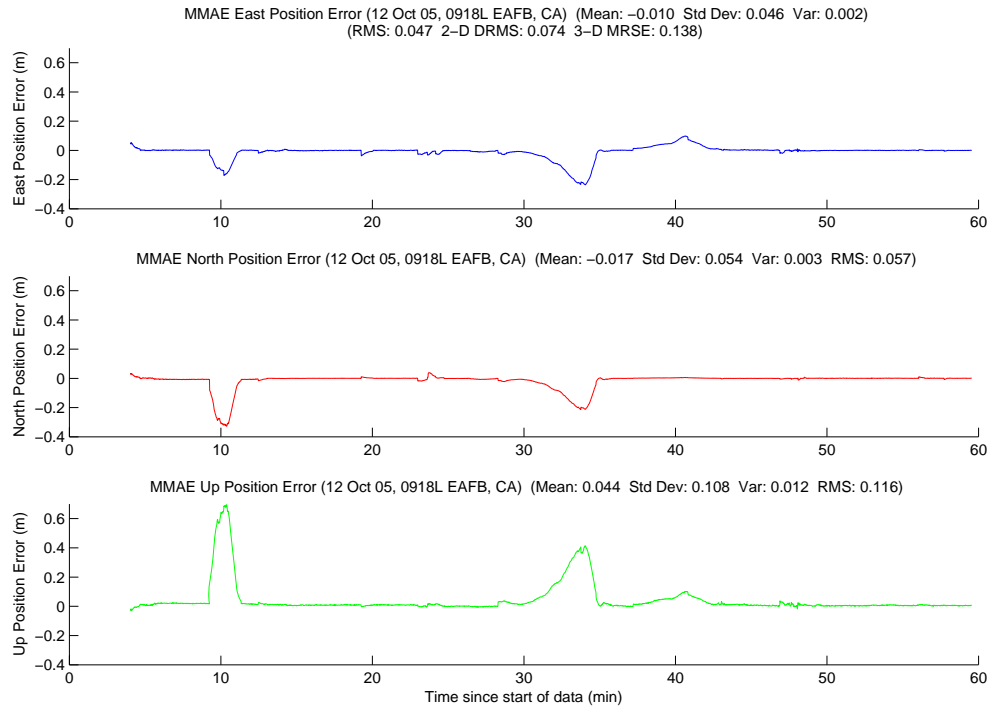


Figure A.256: Case F6.1: MMAE Position Error

Finally, Figure A.257 displays the position error plot for the minimum indicator variable. Since the minimum indicator was “locked” to the true ambiguity set for the entire run, the position errors were incredibly small. Also, due to the extremely short base line distance, the MRSE was within one centimeter for the 60 minute data run! Lastly, Table A.14 gives a summary of each method’s position error.

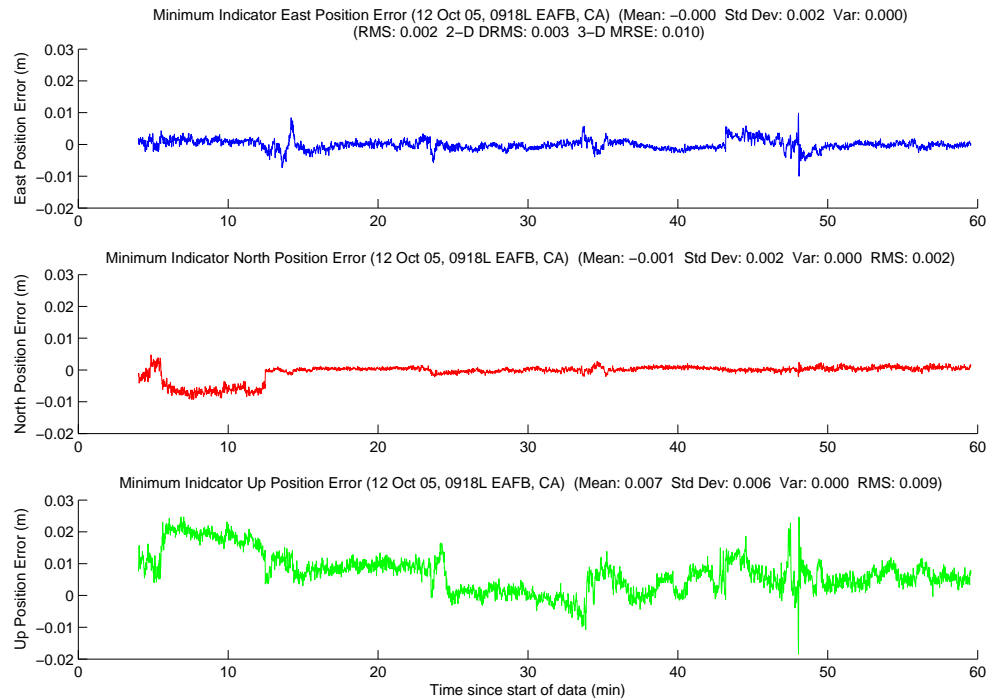


Figure A.257: Case F6.1: Minimum Indicator Position Error

Table A.14: Case F6.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.107	0.048	-0.034	0.058	0.103	0.056	0.136	0.179
MMAE	-0.010	0.046	-0.017	0.054	0.044	0.108	0.074	0.138
Min. Ind.	-0.000	0.002	-0.001	0.002	0.007	0.006	0.003	0.010

A.2.11 Case F6.2: First Flight of Flight Test 6, Second Data Set . The second data set for the first flight started at 1019, one minute after Case 6.1 ended. The run lasted 51 minutes, 50 seconds, and the aircraft ranged from 14.5 to 619.6 meters apart during the data run. Figure A.258 shows the minimum indicator North, East, Down relative positions.

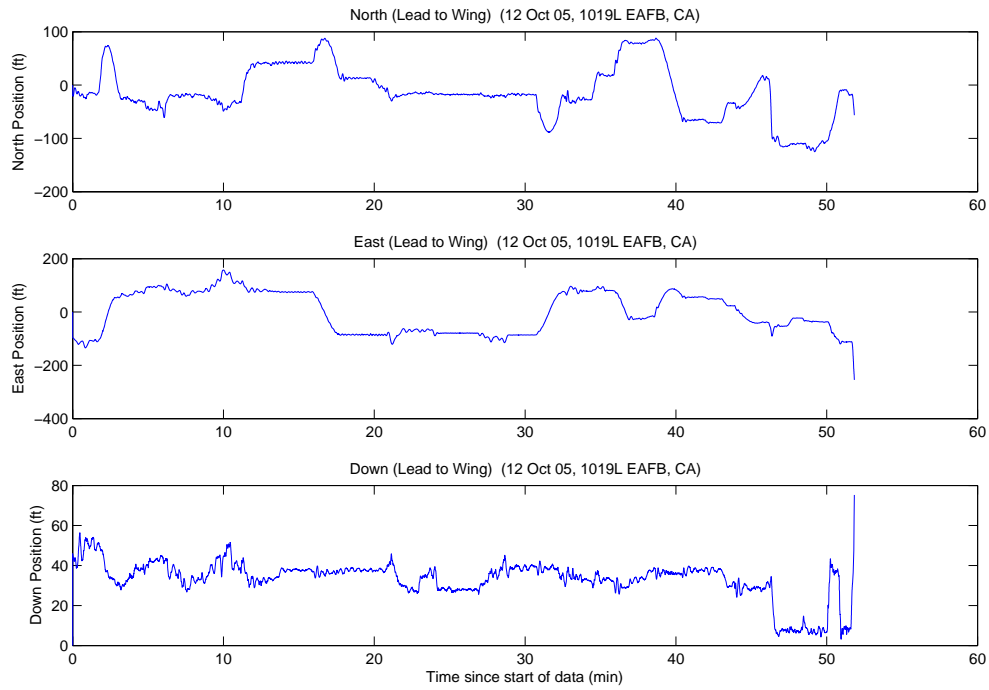


Figure A.258: Case F6.2: North, East, Down Relative Positions

Approximately 30 unique ambiguity sets were returned from the LAMBDA function during the data run. The aircraft were 38.3 meters apart when the system started, with eight visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 2988 of 2988 epochs (100.0 percent).

Figures A.259, A.260, and A.261 show the floating filter errors. Once again, The floating filter performed very well. The floating filter's 1σ estimate of the position error is outside the plot scales in all three figures.

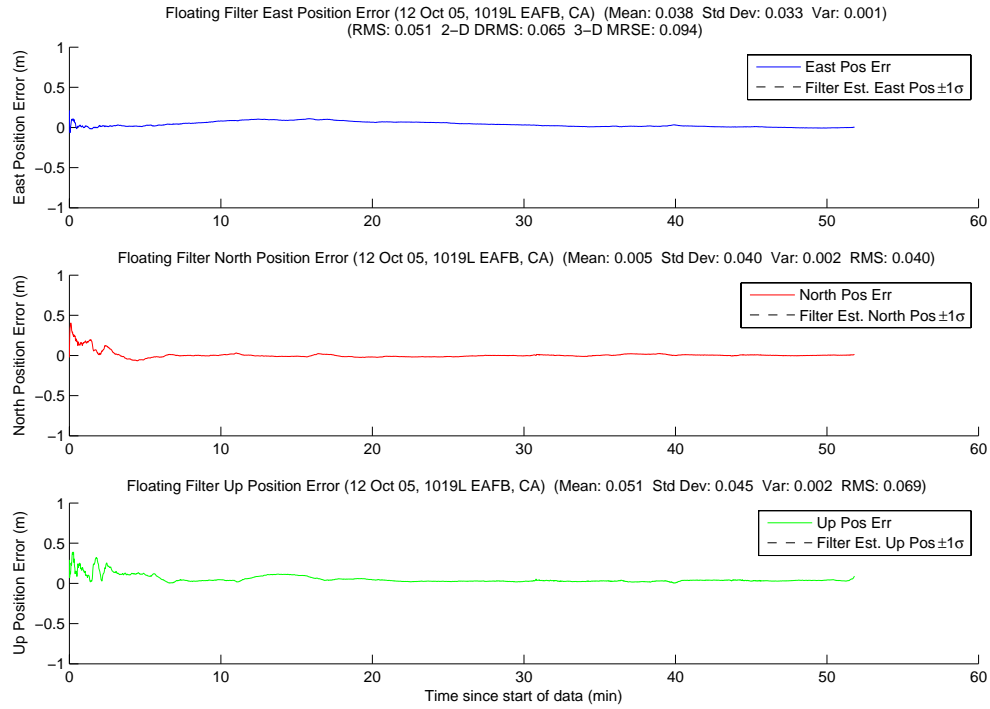


Figure A.259: Case F6.2: Floating Filter Position Errors

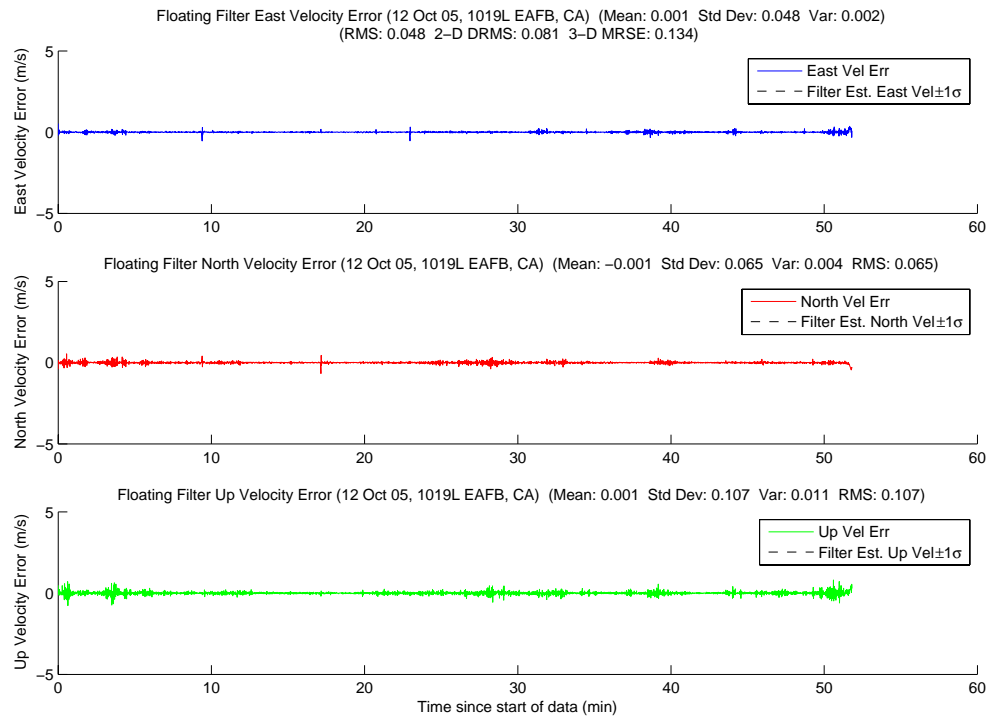


Figure A.260: Case F6.2: Floating Filter Velocity Errors

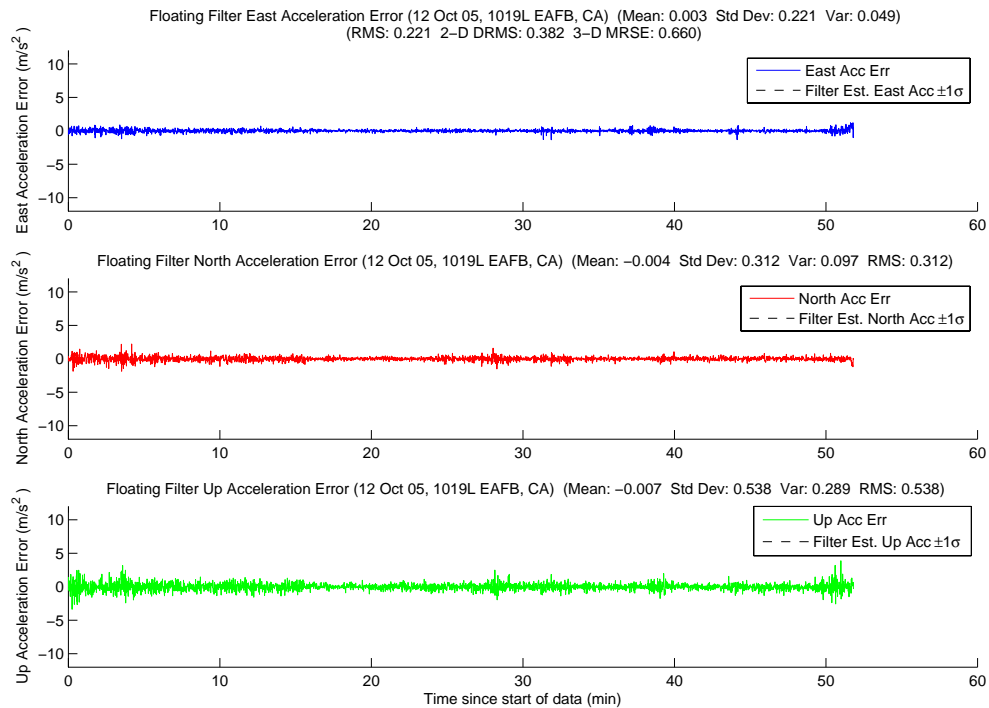


Figure A.261: Case F6.2: Floating Filter Acceleration Errors

Figure A.262 displays PRN 16 measurement information and the floating filter ambiguity estimate. The break in the floating filter's 1σ estimation at 31 minutes was due to masking. The rest of the satellite's plots follow in Figures A.263, A.264, A.265, A.266, A.267, A.268, A.269, and A.270.

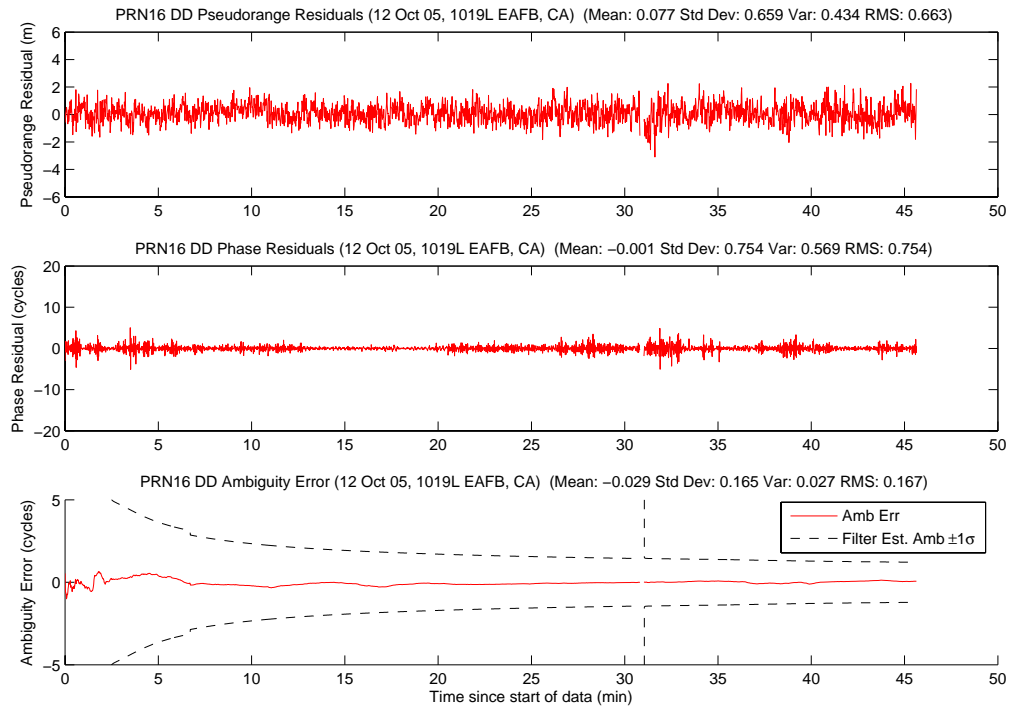


Figure A.262: Case F6.2: Satellite 16 Measurements

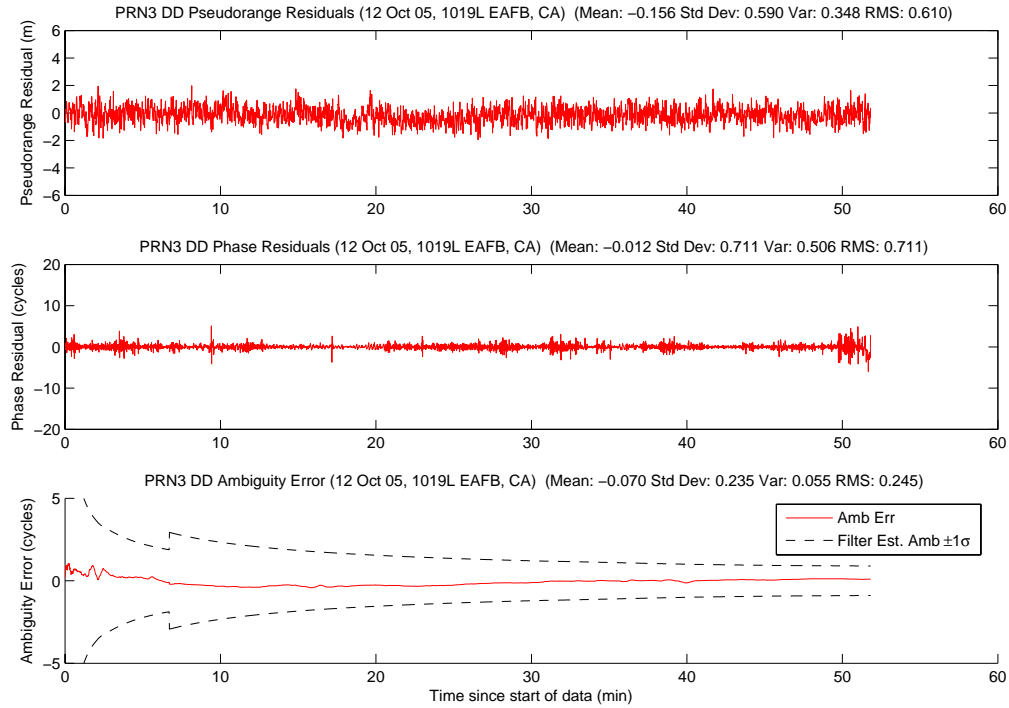


Figure A.263: Case F6.2: Satellite 3 Measurements

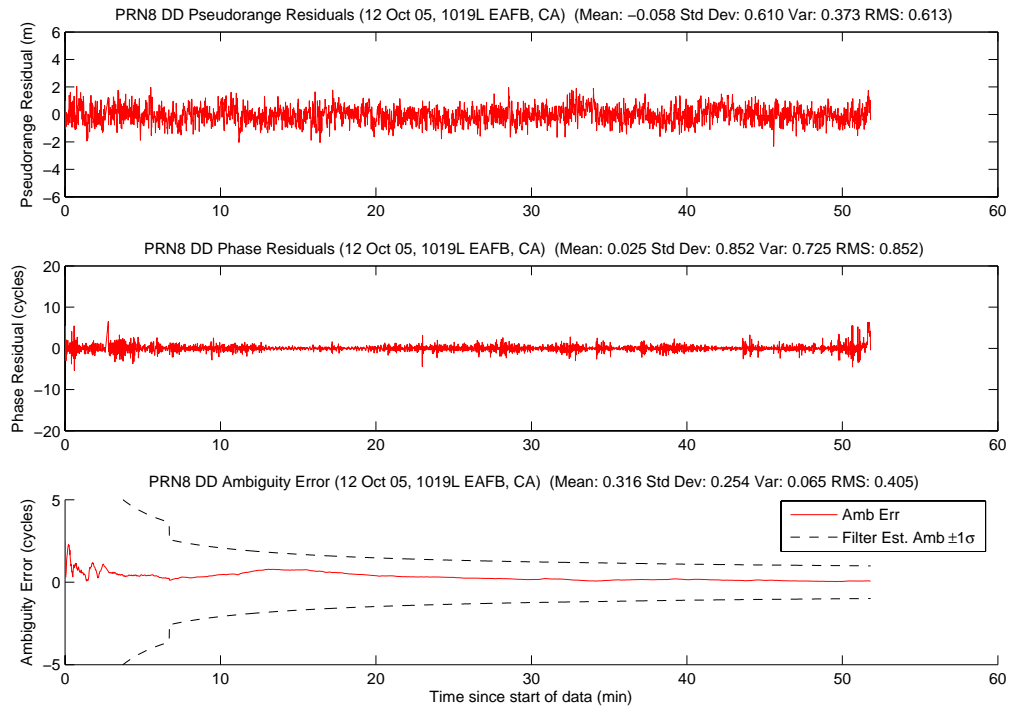


Figure A.264: Case F6.2: Satellite 8 Measurements

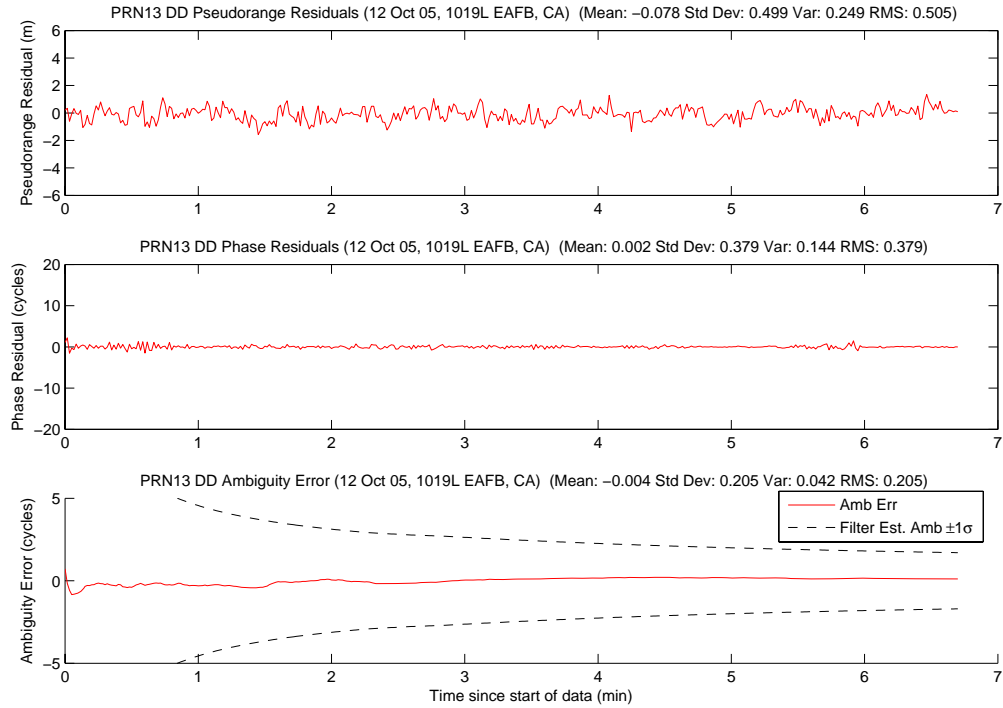


Figure A.265: Case F6.2: Satellite 13 Measurements

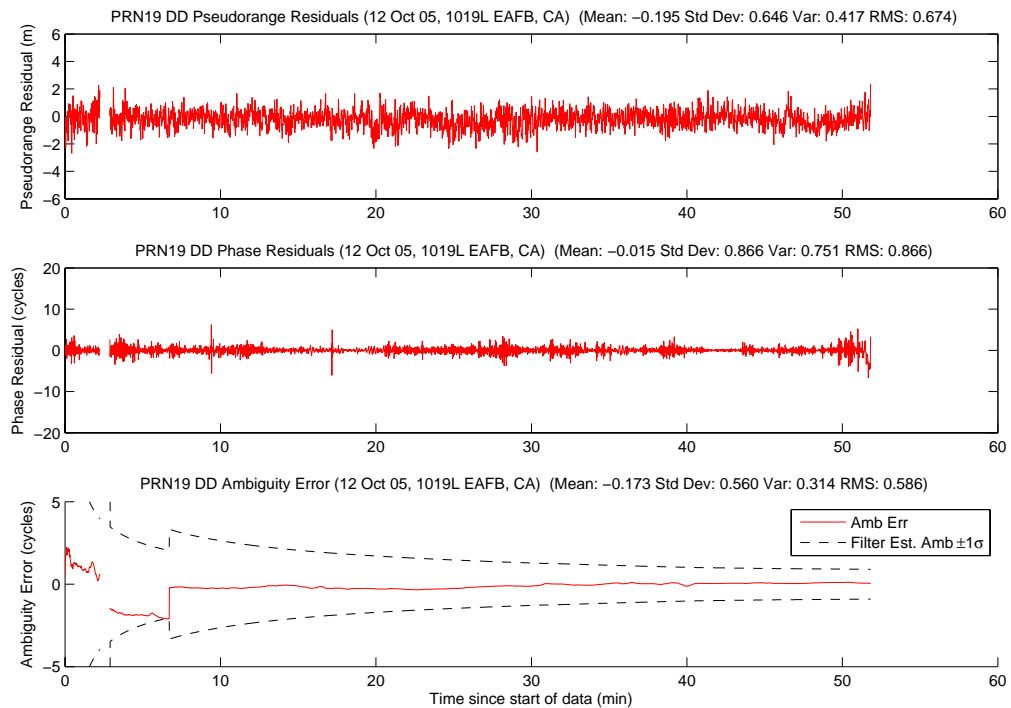


Figure A.266: Case F6.2: Satellite 19 Measurements

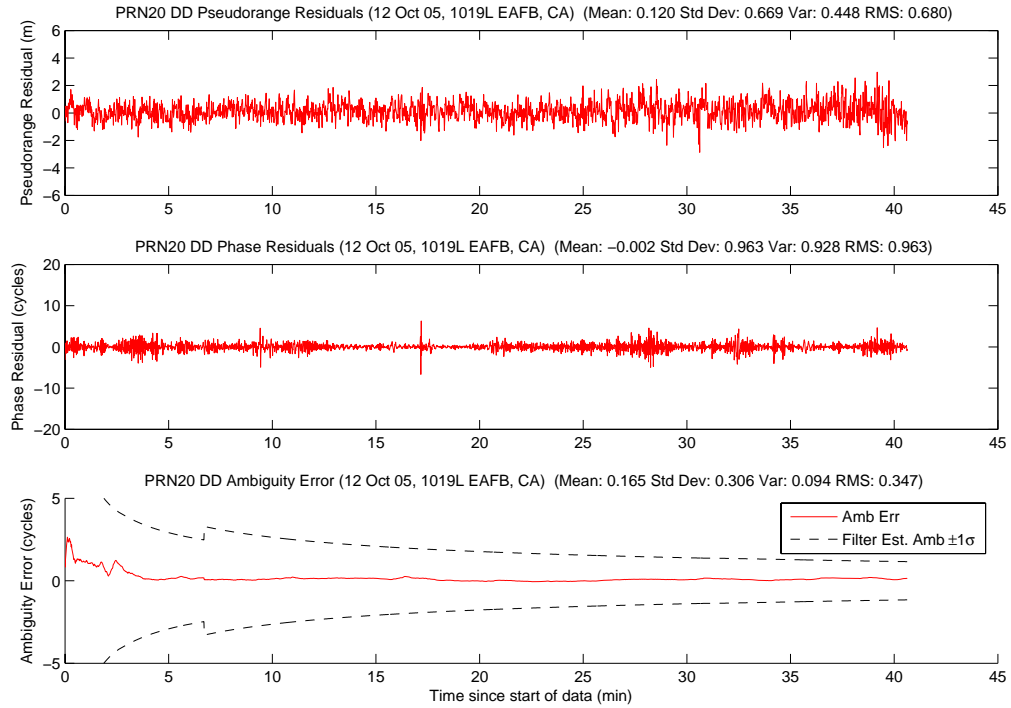


Figure A.267: Case F6.2: Satellite 20 Measurements

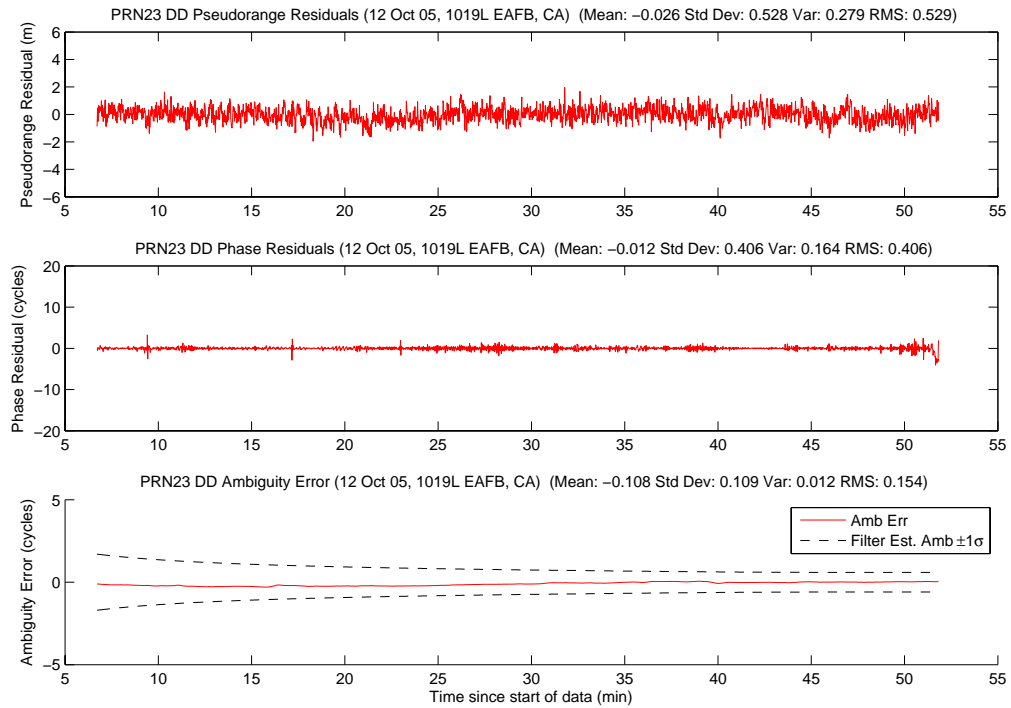


Figure A.268: Case F6.2: Satellite 23 Measurements

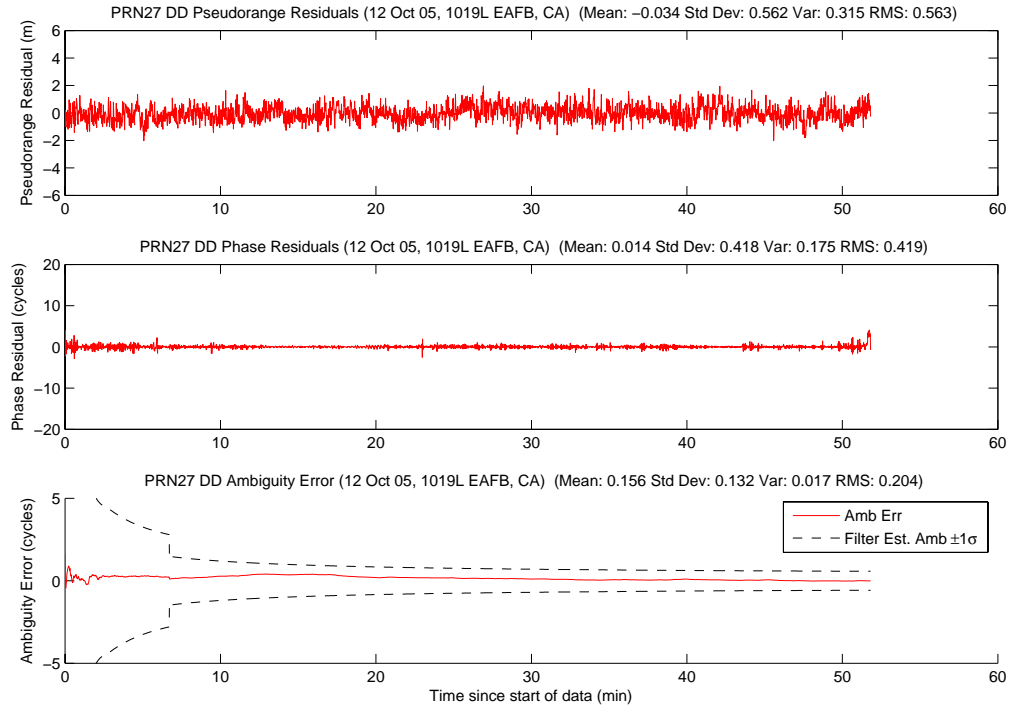


Figure A.269: Case F6.2: Satellite 27 Measurements

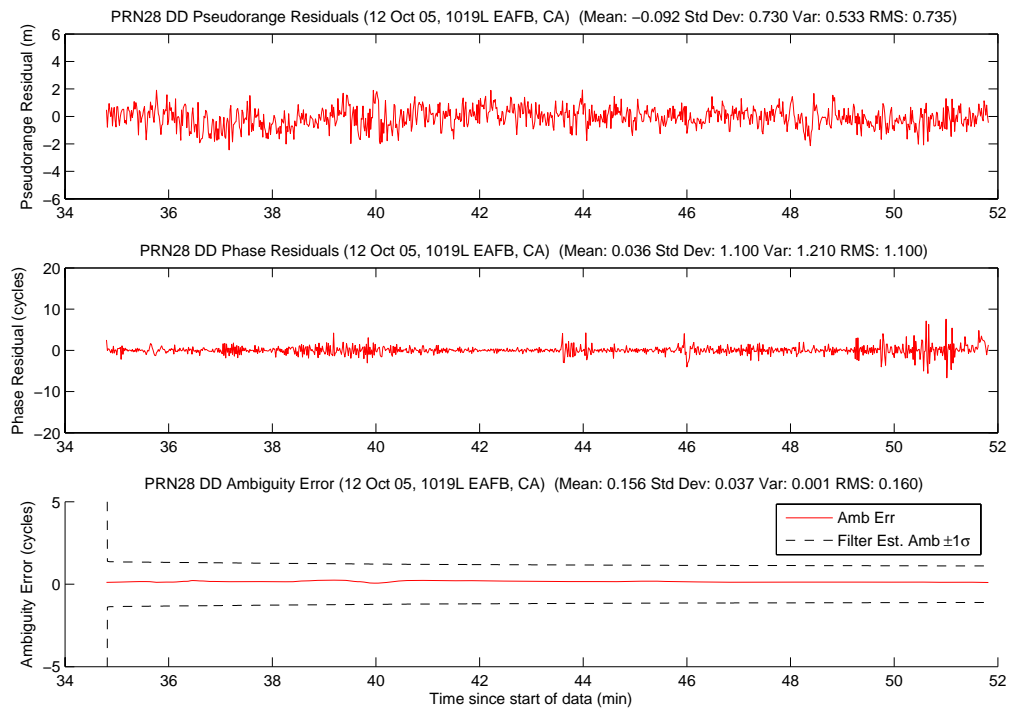


Figure A.270: Case F6.2: Satellite 28 Measurements

Figure A.271 shows the number of visible satellites. All of the satellite dropouts were due to satellite masking. Masking of PRNs 8 and 19 caused the dropouts at three minutes. The dropout at 31 minutes was masking of PRN 16. PRN 28 rose at 35 minutes, and PRNs 20 and 16 set at 40.5 and 45.5 minutes, respectively.

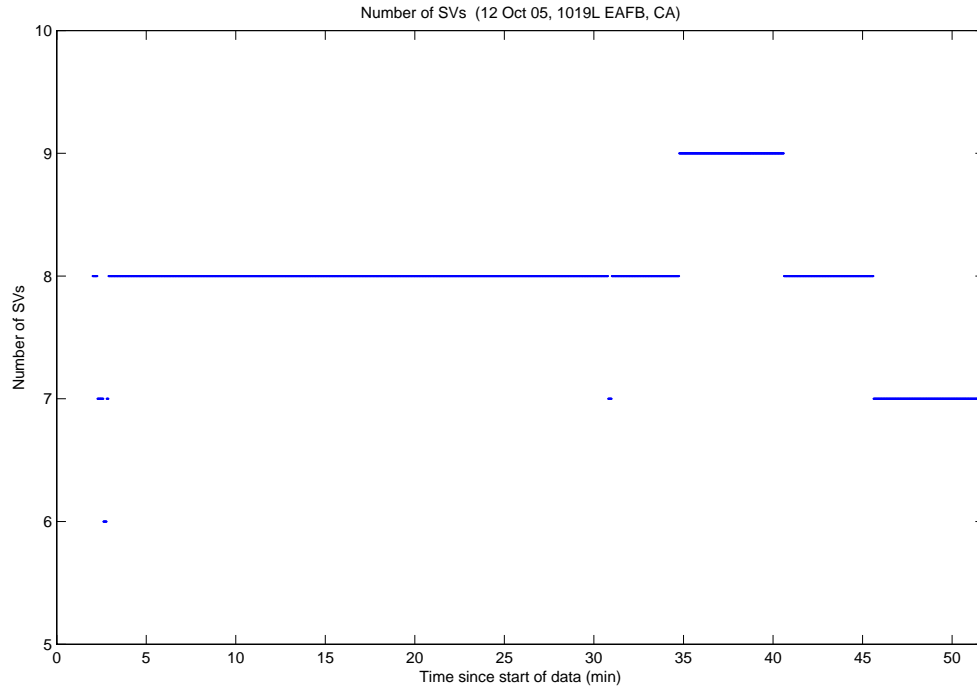


Figure A.271: Case F6.2: Visible Satellites

Figure A.272 shows all of the candidate ambiguity set's SOS residuals. Case 6.2 residual plot was “cleaner” than the previous case. The true ambiguity set stood out much clearer because of the information garnered from just two extra satellites (eight visible satellites compared to just six in Case 6.1) made a large difference. The only time when an erroneous ambiguity set's residuals were below the true ambiguity set was at 15.5 minutes. Figure A.273 displays the MMAE conditional probability. There was only one large dip in the true ambiguity set's probability at approximately 15 minutes, which corresponded to the residual plot discussed previously.

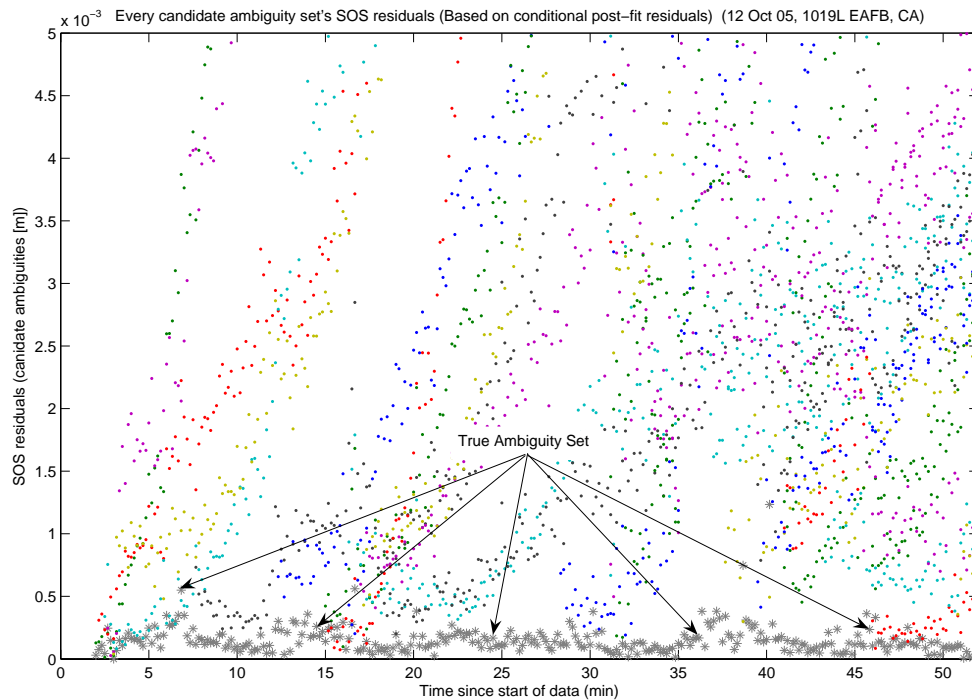


Figure A.272: Case F6.2: SOS Residuals

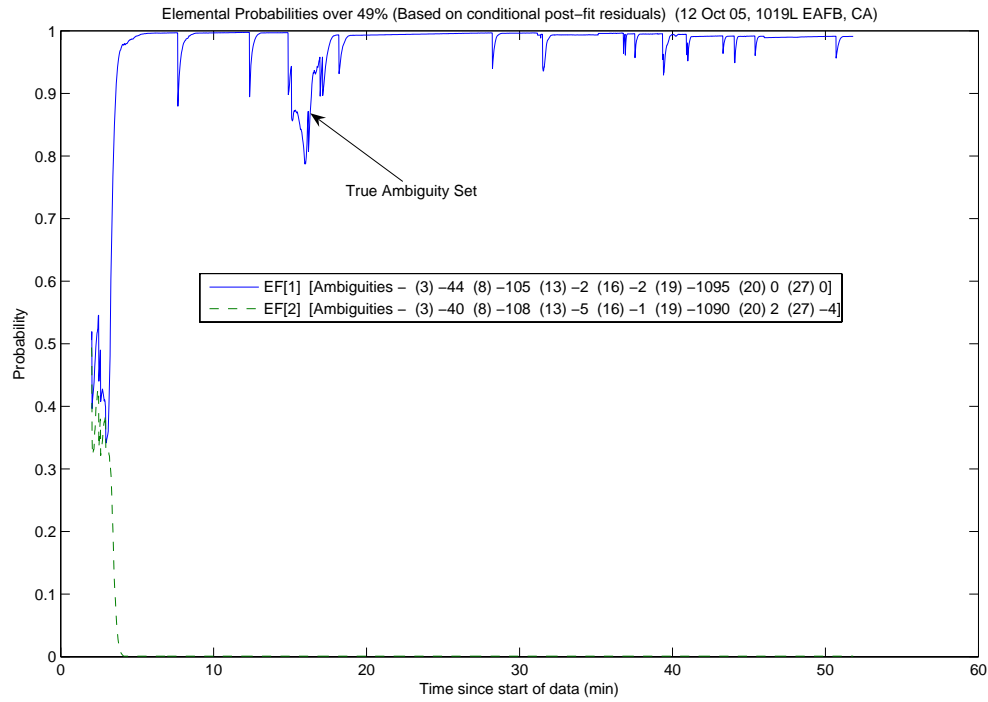


Figure A.273: Case F6.2: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.274 is EF[1], and Figure A.275 is EF[2].

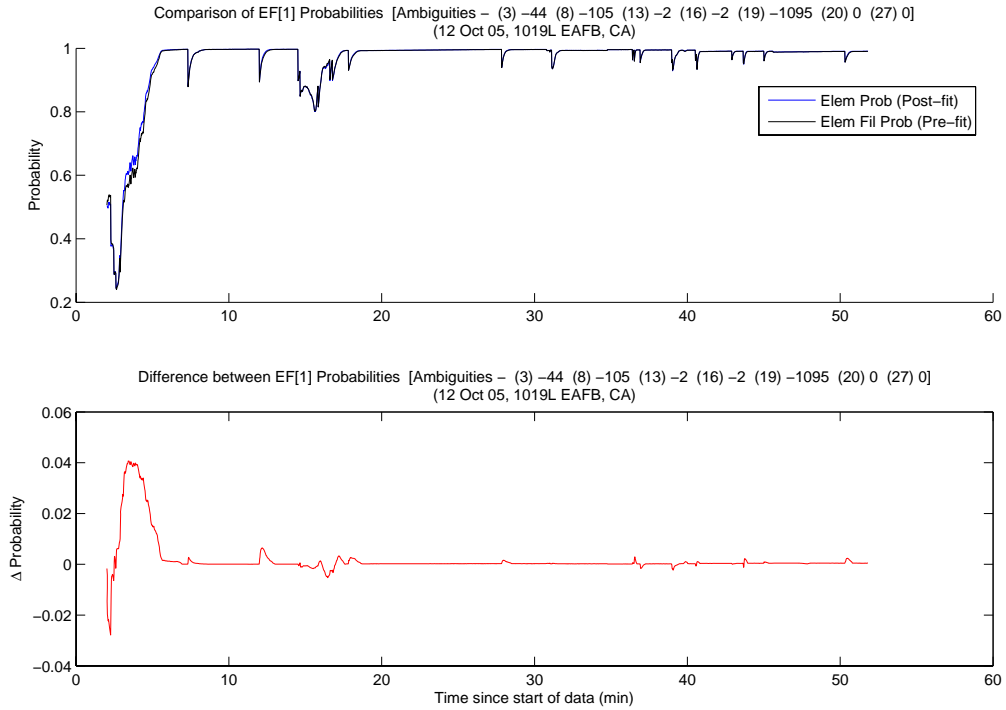


Figure A.274: Case F6.2: EF[1] Probability Comparison

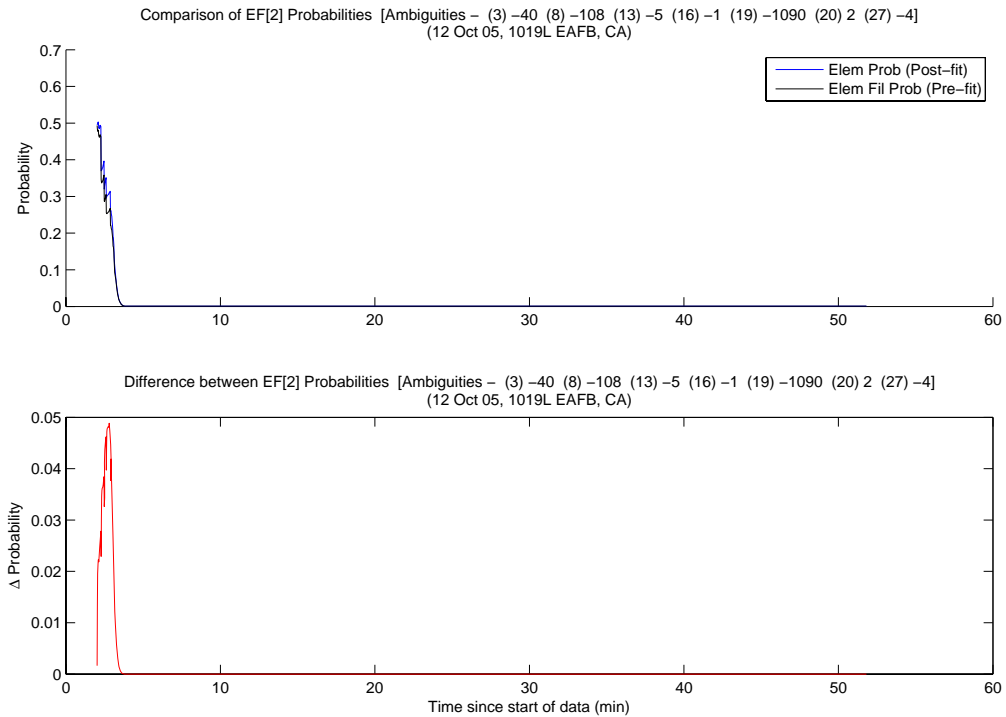


Figure A.275: Case F6.2: EF[2] Probability Comparison

The MMAE position error is shown in Figure A.276. The error in the position solution, due to the probability drop at 15 minutes, is evident in the plot. Otherwise, the MMAE position solution was generally accurate.

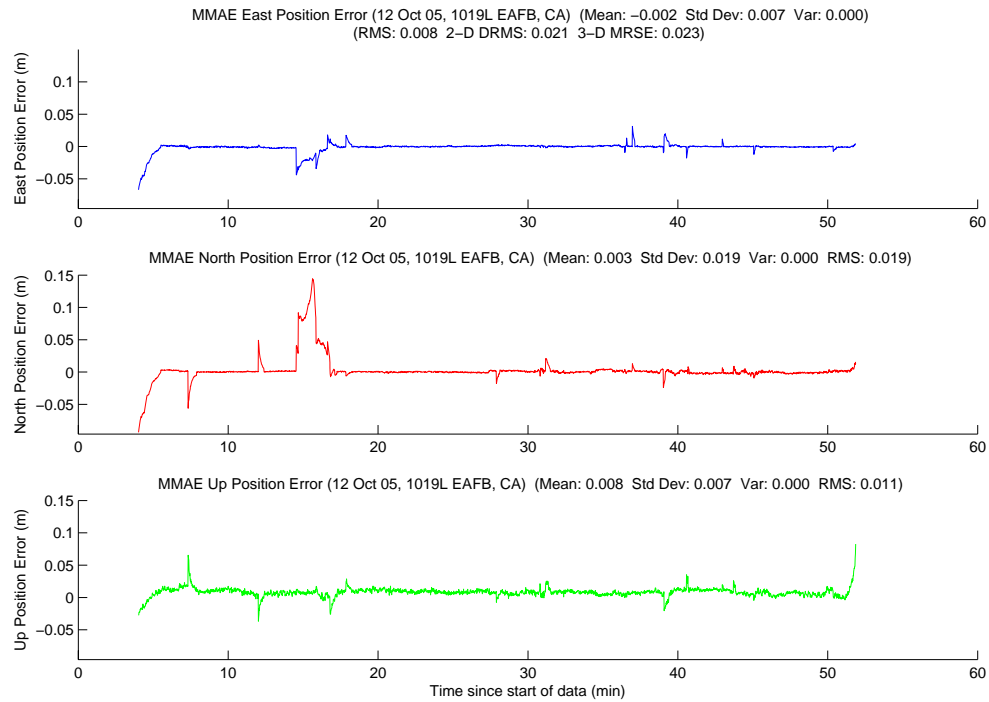


Figure A.276: Case F6.2: MMAE Position Error

Finally, Figure A.277 displays the position error plot for the minimum indicator variable. Since the minimum indicator was “locked” to the true ambiguity set for the entire run, the position errors were incredibly small. Once more, the MRSE was within one centimeter for the 50 minute data run. Lastly, Table A.15 gives a summary of each method’s position error.

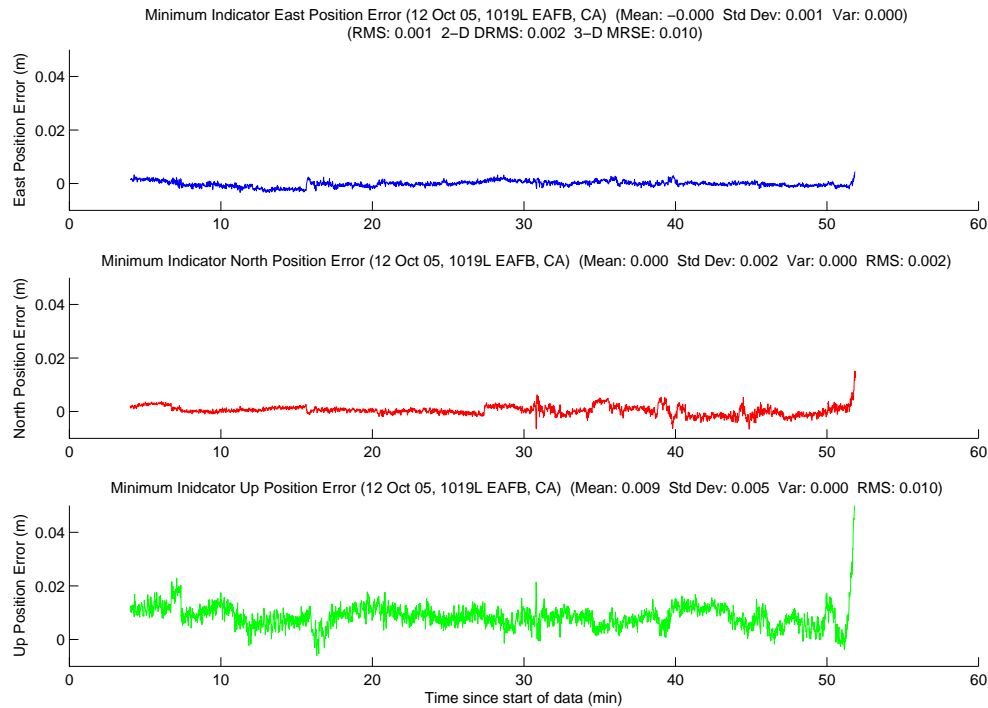


Figure A.277: Case F6.2: Minimum Indicator Position Error

Table A.15: Case F6.2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.038	0.033	0.005	0.040	0.051	0.045	0.065	0.094
MMAE	-0.002	0.007	0.003	0.019	0.008	0.007	0.021	0.023
Min. Ind.	-0.000	0.001	0.000	0.002	0.009	0.005	0.002	0.010

A.2.12 Case F6.3: Second Flight of Flight Test 6, First Data Set . The second flight's first data run, on 12 October 2005, started at 1440L and lasted 59 minutes, 33 seconds. The aircraft ranged from 12.4 to 228.5 meters apart during the data run. Figure A.278 shows the minimum indicator North, East, Down relative positions.

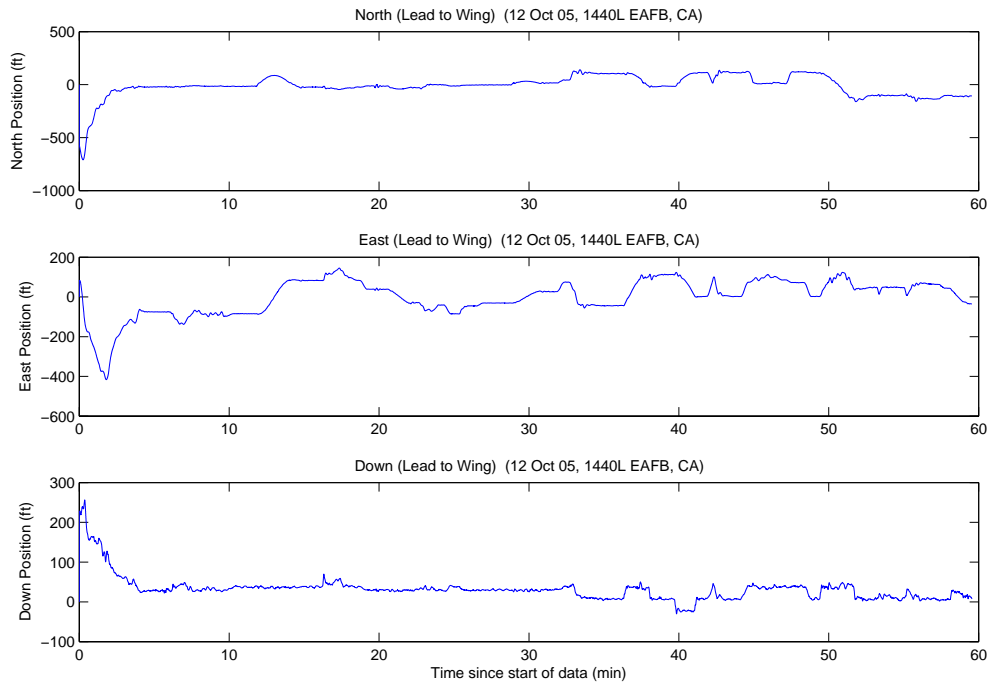


Figure A.278: Case F6.3: North, East, Down Relative Positions

A total of 33 unique ambiguity sets were returned from the LAMBDA function during the data run. The aircraft were 188.7 meters apart when the system started, with nine visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 3452 of 3452 epochs (100.0 percent), with a maximum ratio of 243.6.

Figures A.279, A.280, and A.281 show the floating filter errors. The floating filter's 1σ estimate of the position error is outside the plot scales in all three figures.

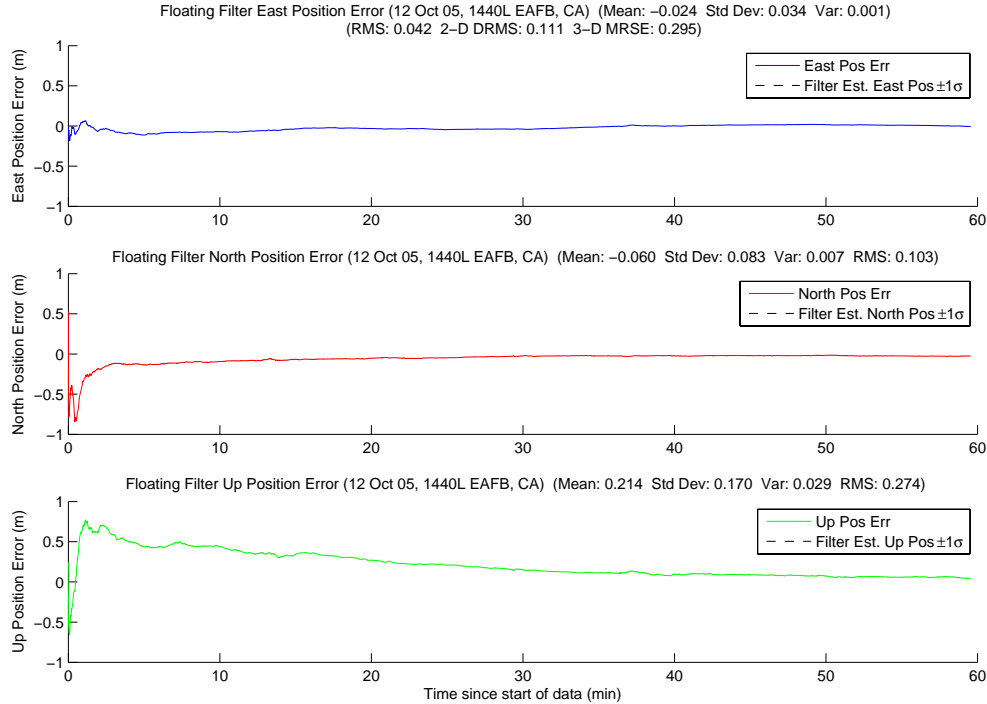


Figure A.279: Case F6.3: Floating Filter Position Errors

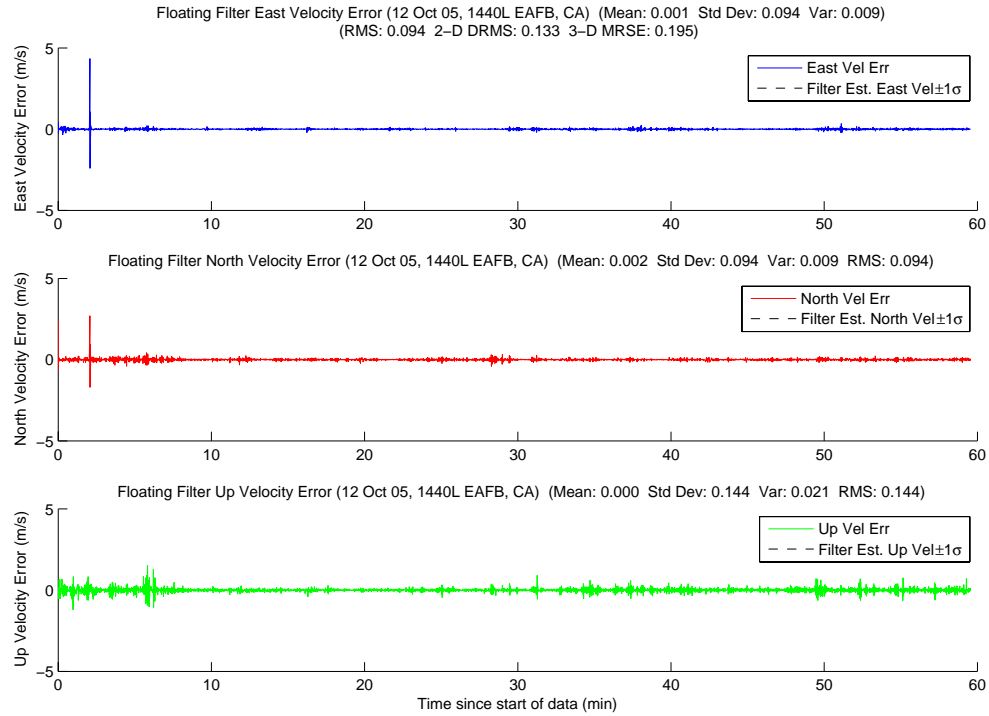


Figure A.280: Case F6.3: Floating Filter Velocity Errors

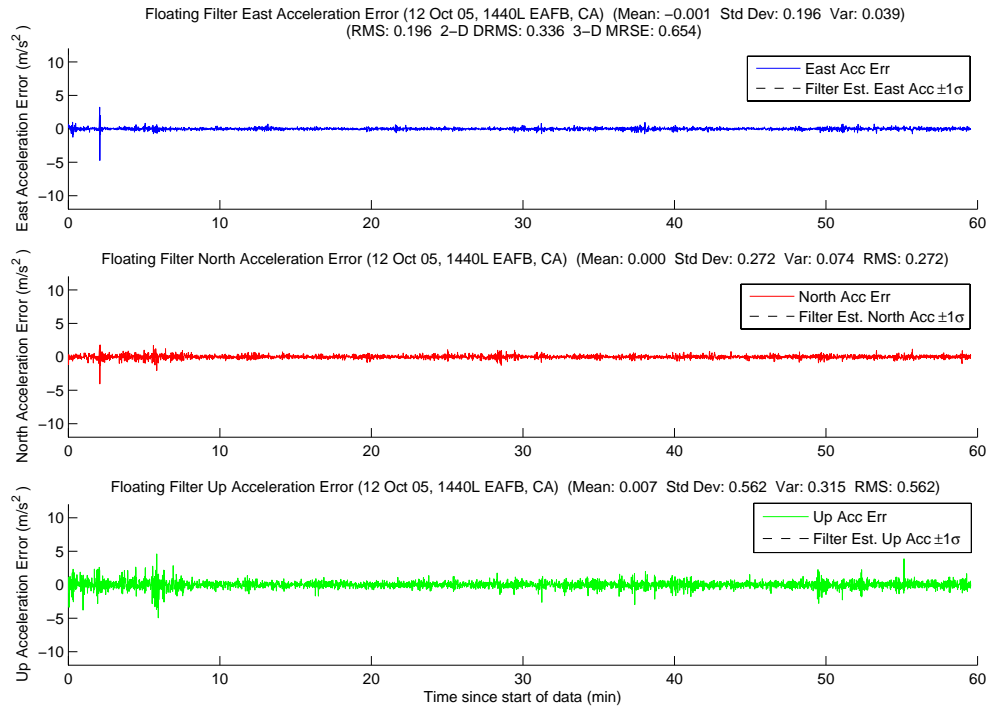


Figure A.281: Case F6.3: Floating Filter Acceleration Errors

Figure A.282 displays PRN 17 measurement information and the floating filter ambiguity estimate. The jump in the pseudorange residuals and the floating filter's 1σ estimation was due to a change in the base satellite from PRN 31 to PRN 7. The rest of the satellite's plots follow in Figures A.283, A.284, A.285, A.286, A.287, A.288, A.289, A.290, A.291, A.292, and A.293.

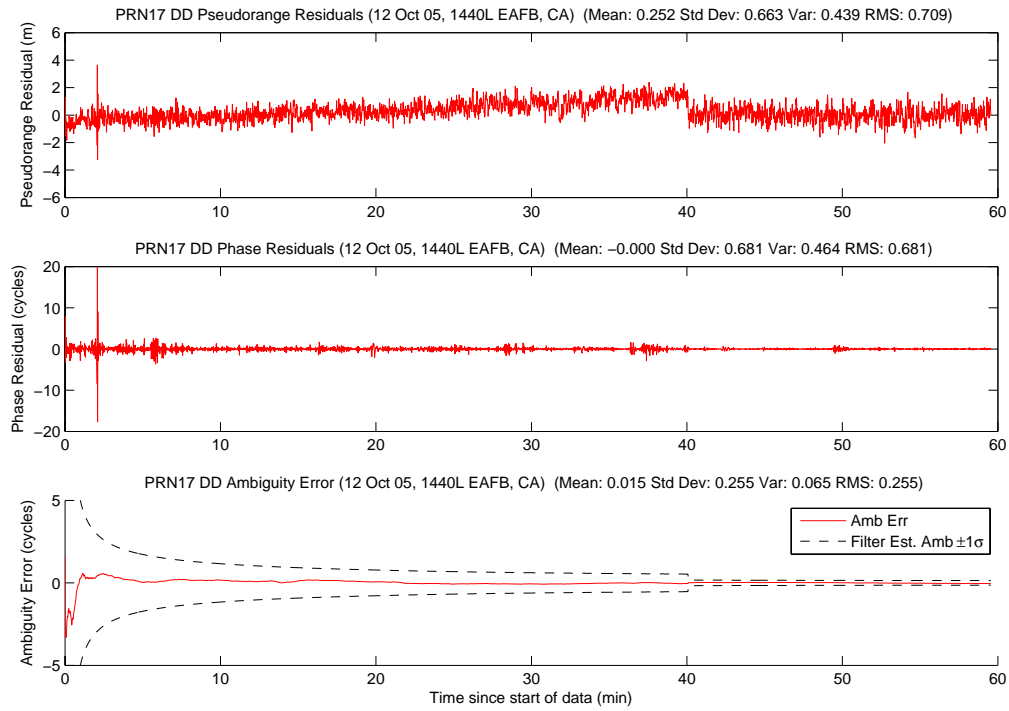


Figure A.282: Case F6.3: Satellite 17 Measurements

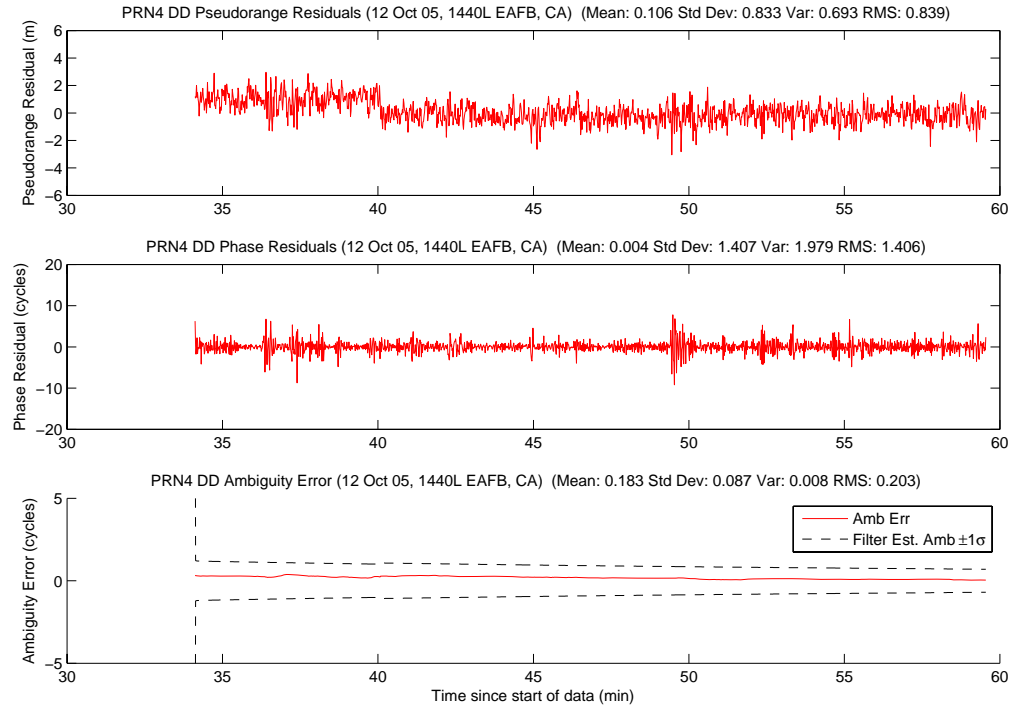


Figure A.283: Case F6.3: Satellite 4 Measurements

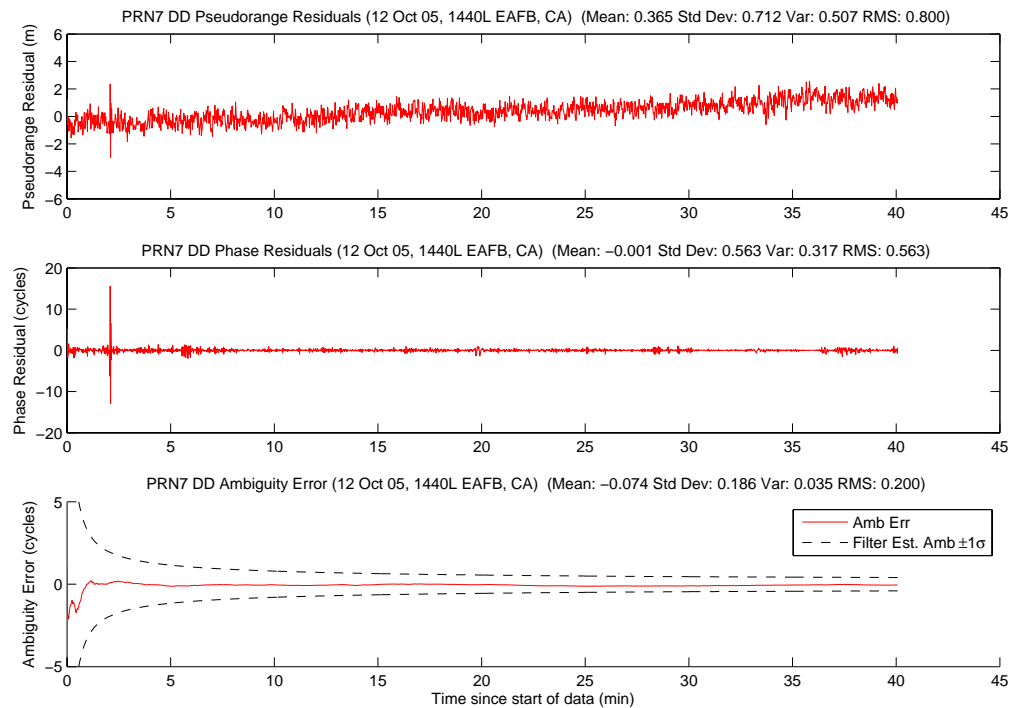


Figure A.284: Case F6.3: Satellite 7 Measurements

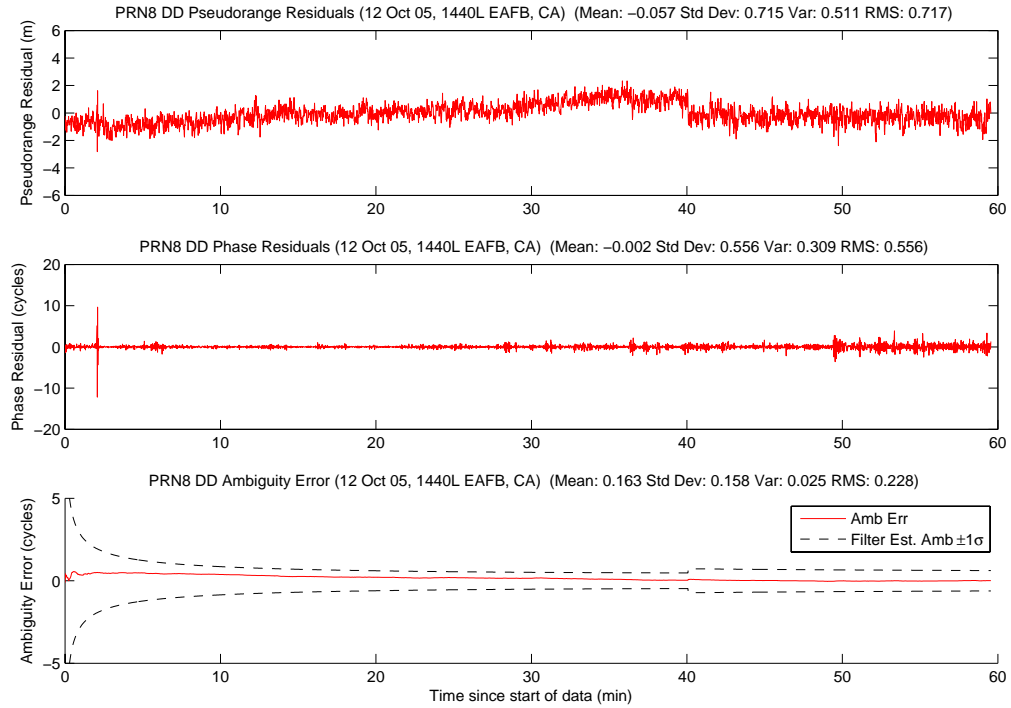


Figure A.285: Case F6.3: Satellite 8 Measurements

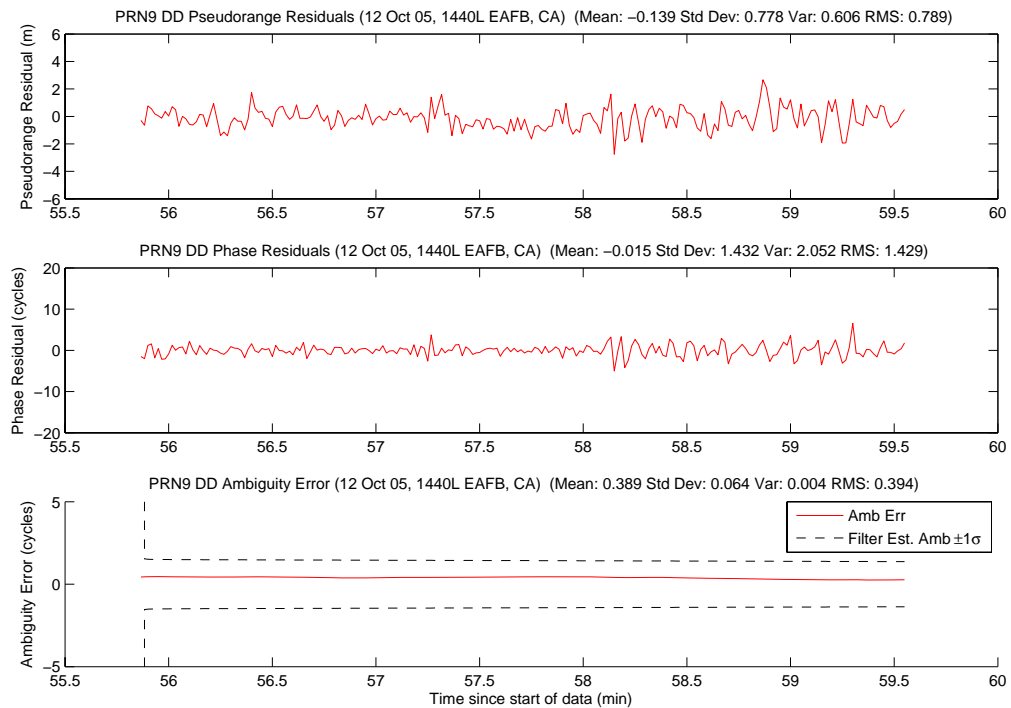


Figure A.286: Case F6.3: Satellite 9 Measurements

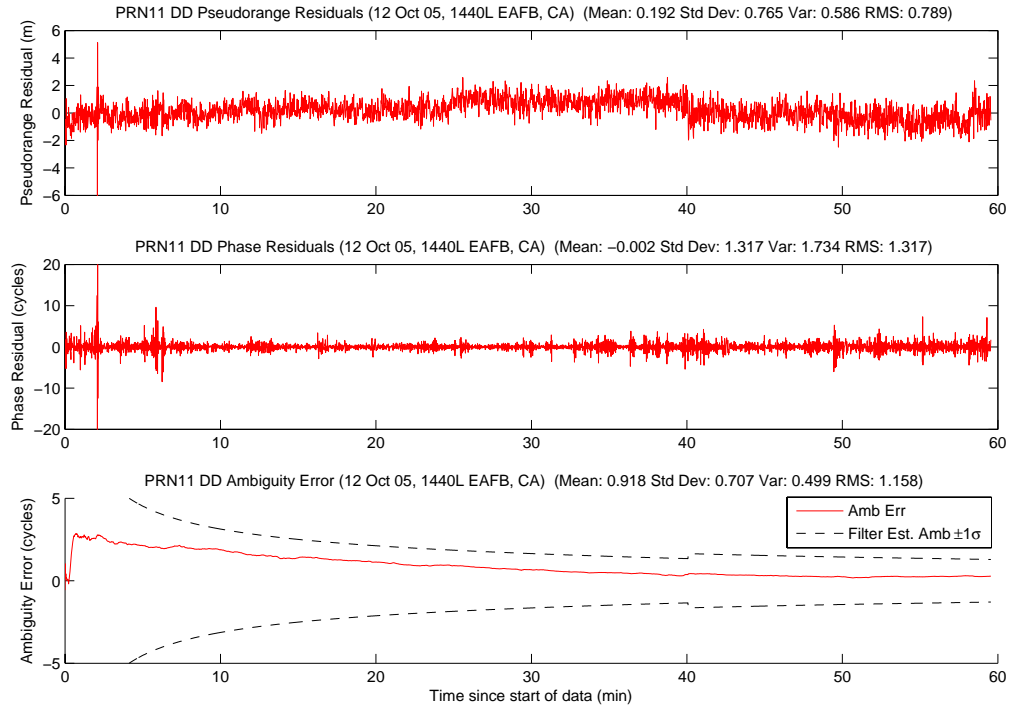


Figure A.287: Case F6.3: Satellite 11 Measurements

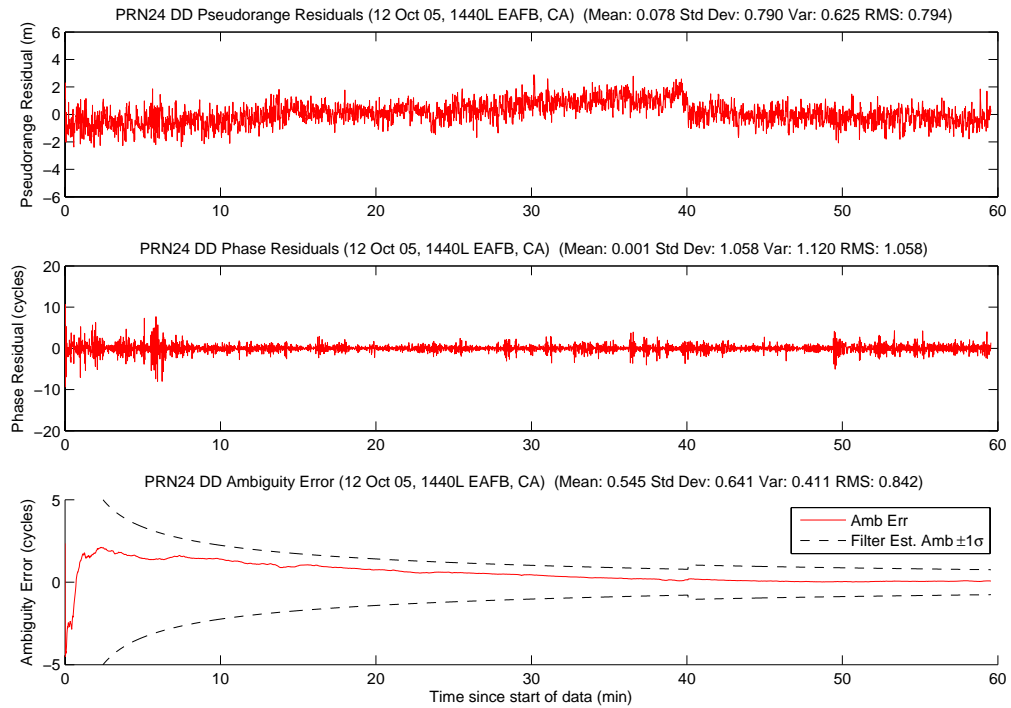


Figure A.288: Case F6.3: Satellite 24 Measurements

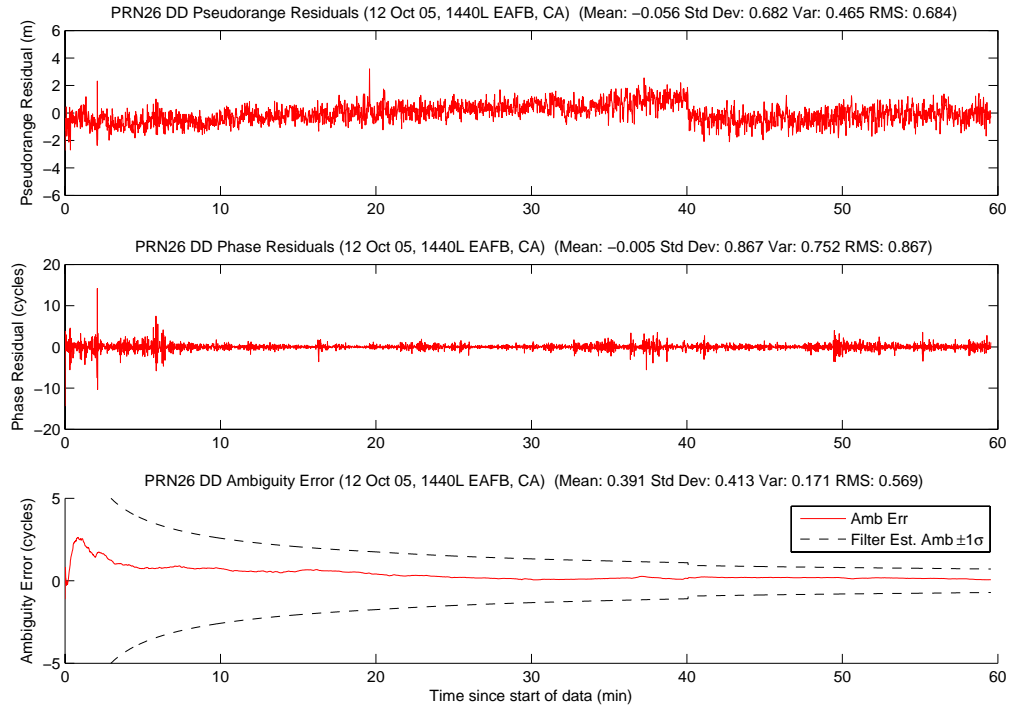


Figure A.289: Case F6.3: Satellite 26 Measurements

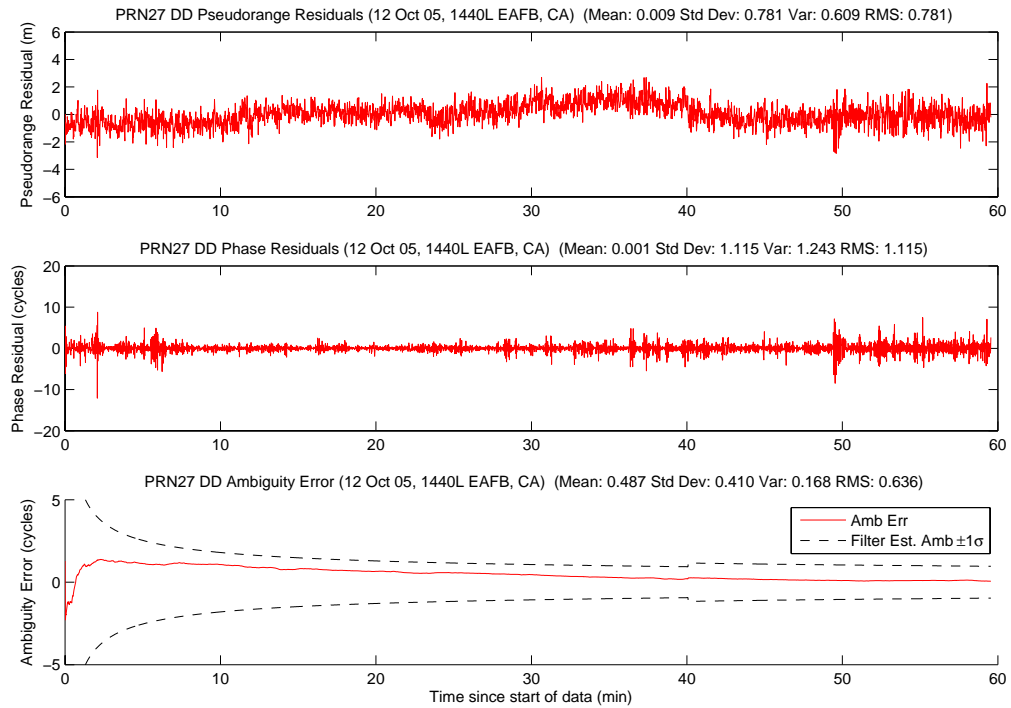


Figure A.290: Case F6.3: Satellite 27 Measurements

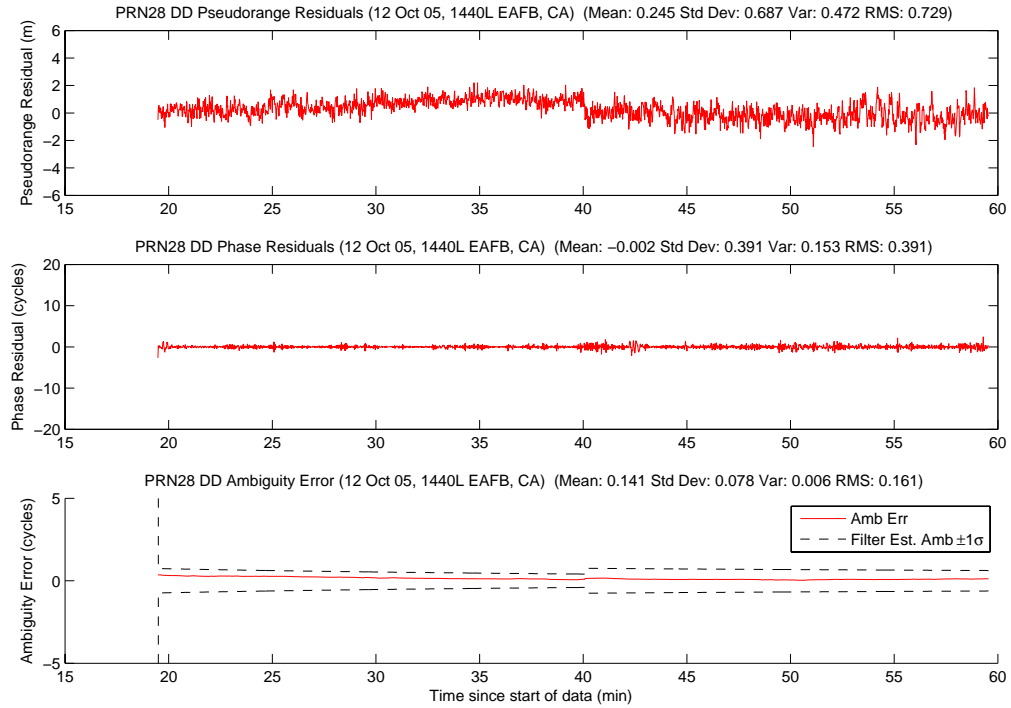


Figure A.291: Case F6.3: Satellite 28 Measurements

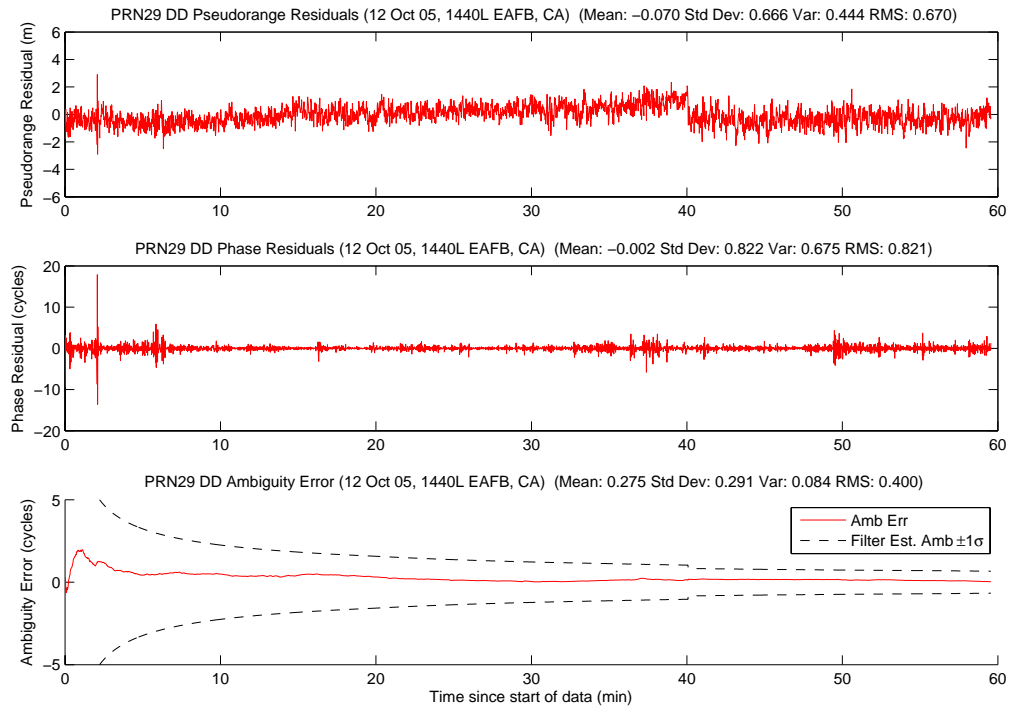


Figure A.292: Case F6.3: Satellite 29 Measurements

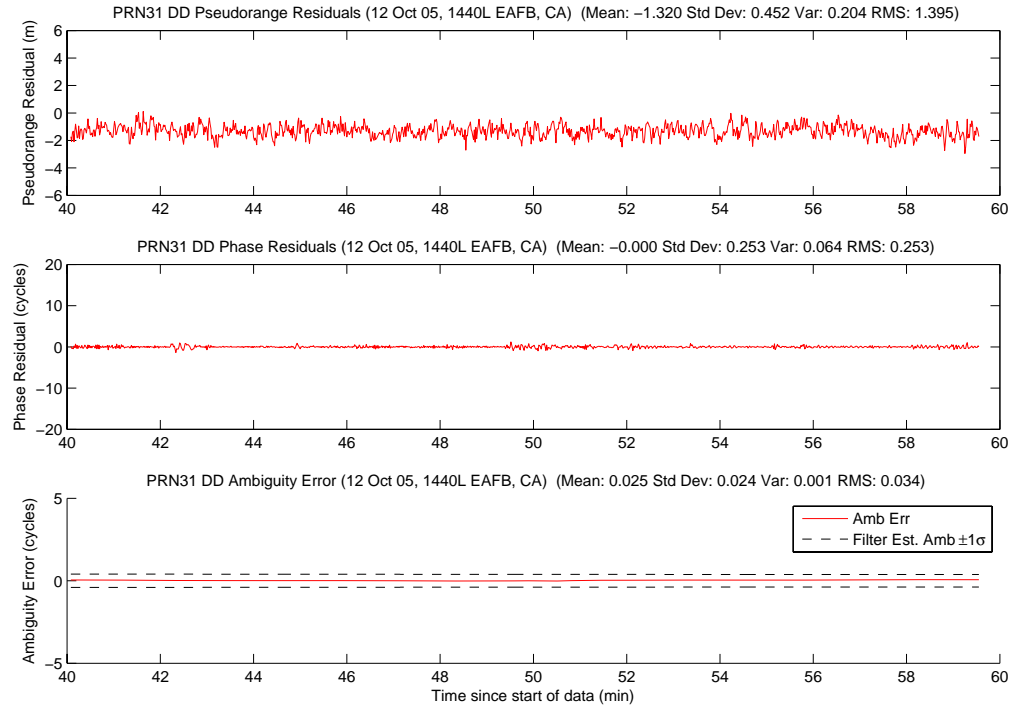


Figure A.293: Case F6.3: Satellite 31 Measurements

Figure A.294 shows the number of visible satellites. There were no satellite dropouts during the run. In fact, no satellites set during the run, but three other satellites were gained. PRNs 28, 4, and 9 rose at 19.5, 34, and 56 minutes respectively.

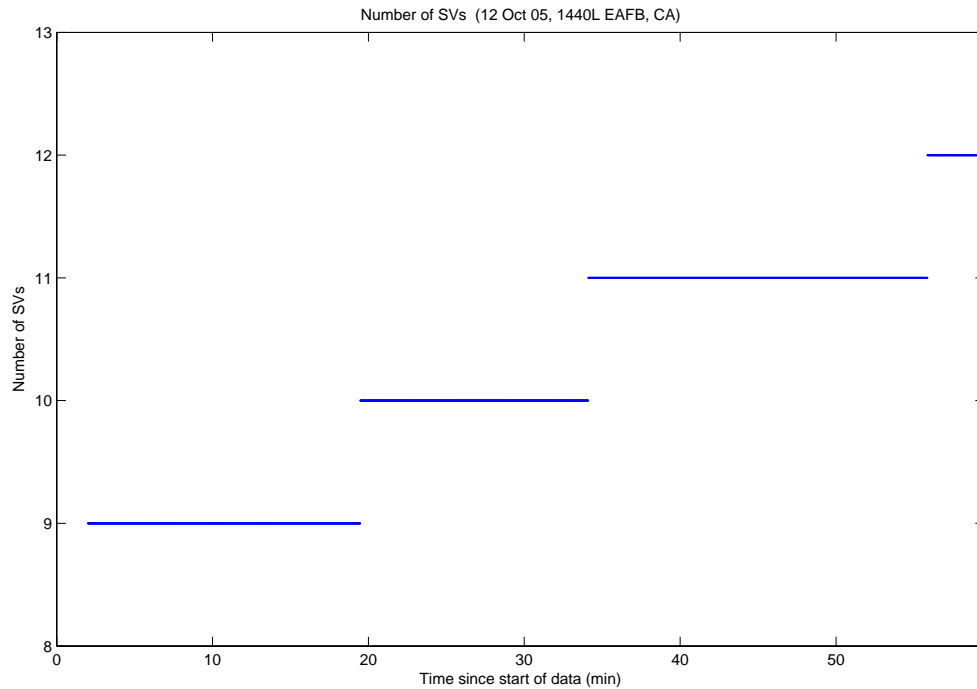


Figure A.294: Case F6.3: Visible Satellites

Figure A.295 shows all of the candidate ambiguity set's SOS residuals. The residual plot is similar to Cases 5.1 and 5.2; once more, showing the power of the number of satellites and their associated geometry. Figure A.296 displays the MMAE conditional probability. After a few dips, the true ambiguity set's probability remains above 90 percent for the entire run.

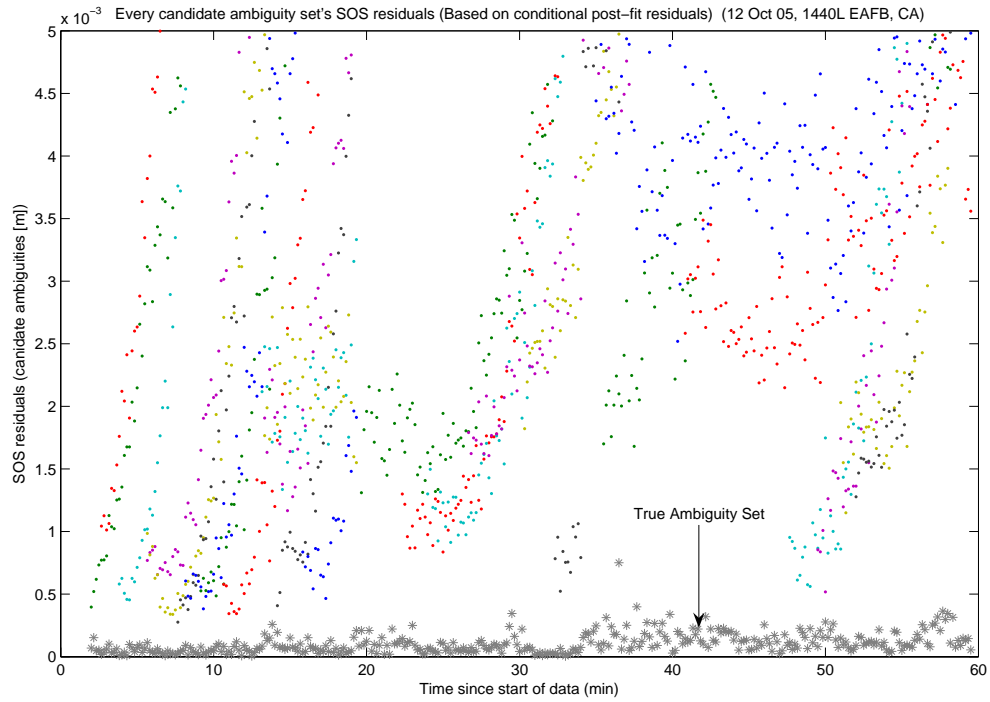


Figure A.295: Case F6.3: SOS Residuals

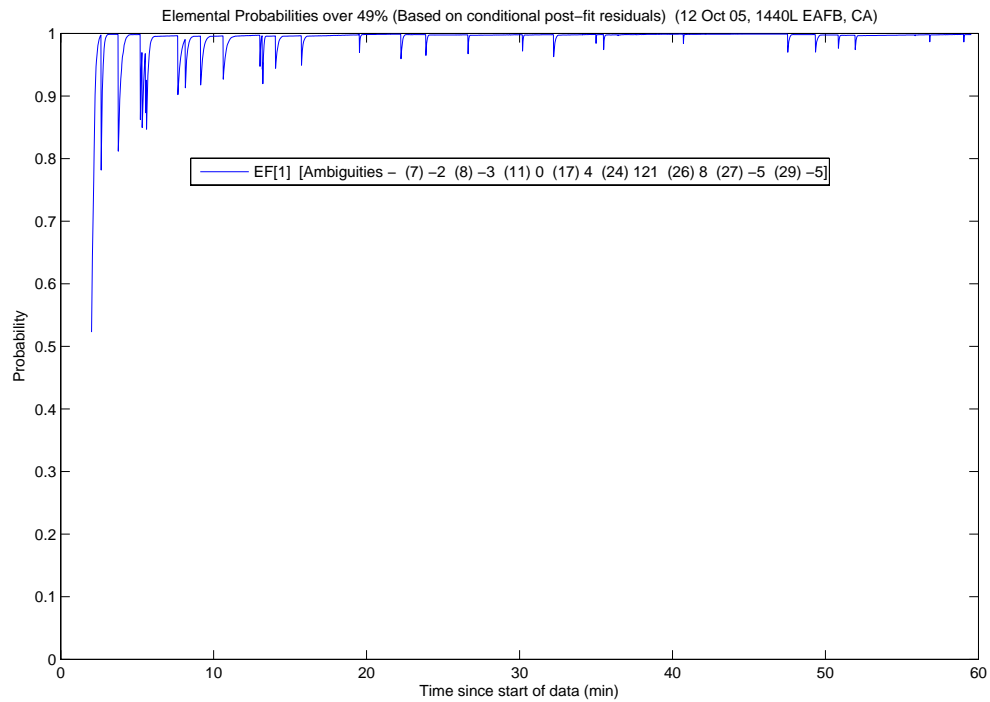


Figure A.296: Case F6.3: Conditional Probabilities for Selected MMAE Elemental Filters

The following figure shows the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.297 is EF[1].

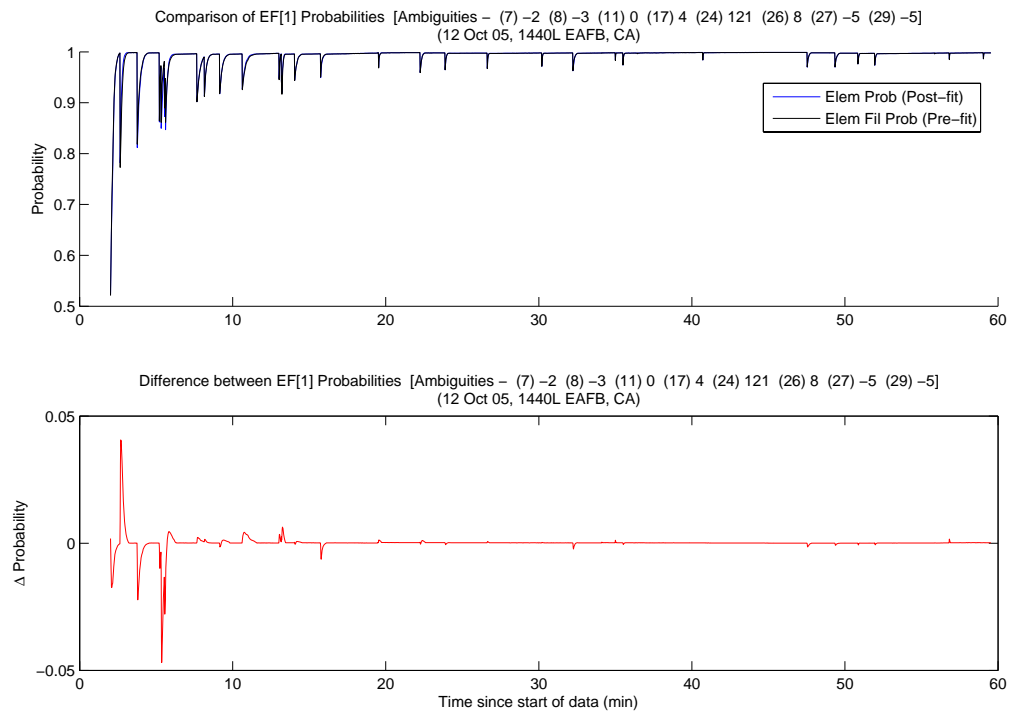


Figure A.297: Case F6.3: EF[1] Probability Comparison

The MMAE position error is shown in Figure A.298. The few dips in the true ambiguity's probability are manifested in position error at the beginning of the data run. After that, MMAE position solution remained accurate.

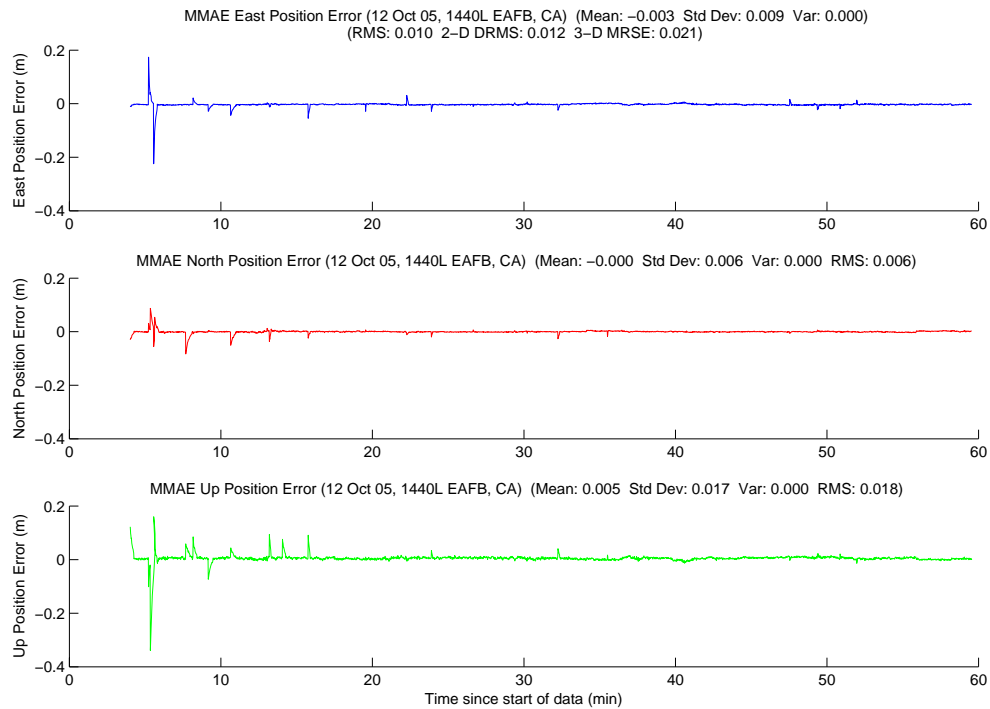


Figure A.298: Case F6.3: MMAE Position Error

Finally, Figure A.299 displays the position error plot for the minimum indicator variable. For the first time, the average minimum indicator position error was within millimeters! Lastly, Table A.16 gives a summary of each method's position error.

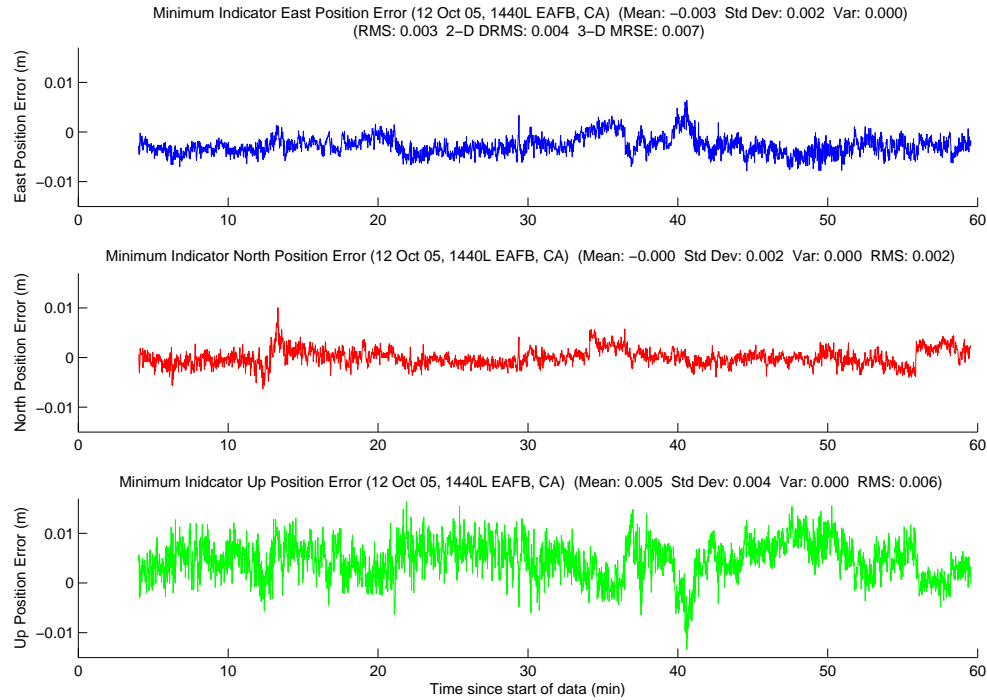


Figure A.299: Case F6.3: Minimum Indicator Position Error

Table A.16: Case F6.3: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.024	0.034	-0.060	0.083	0.214	0.170	0.111	0.295
MMAE	-0.003	0.009	-0.000	0.006	0.005	0.017	0.012	0.021
Min. Ind.	-0.003	0.002	-0.000	0.002	0.005	0.004	0.004	0.007

A.2.13 Case F6.4: Second Flight of Flight Test 6, Second Data Set . The final data run for Case F6 started at 1540L, immediately after Case F6.3 ended. The run lasted 30 minutes, 55 seconds and the aircraft ranged from 11.6 to 116.5 meters apart. Figure A.300 shows the minimum indicator North, East, Down relative positions.

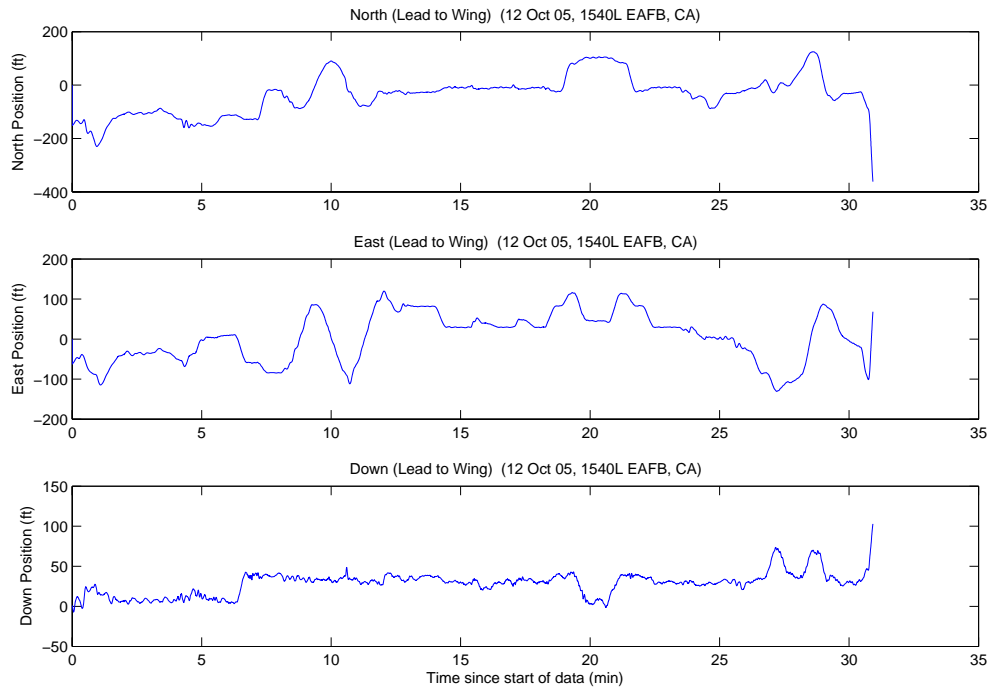


Figure A.300: Case F6.4: North, East, Down Relative Positions

A total of 20 unique ambiguity sets were returned from the LAMBDA function during the data run. The aircraft were 49.8 meters apart when the system started, with twelve visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 1734 of 1734 epochs (100.0 percent).

Figures A.301, A.302, and A.303 show the floating filter errors. The floating filter's 1σ estimate of the position error is outside the plot scales in all three figures.

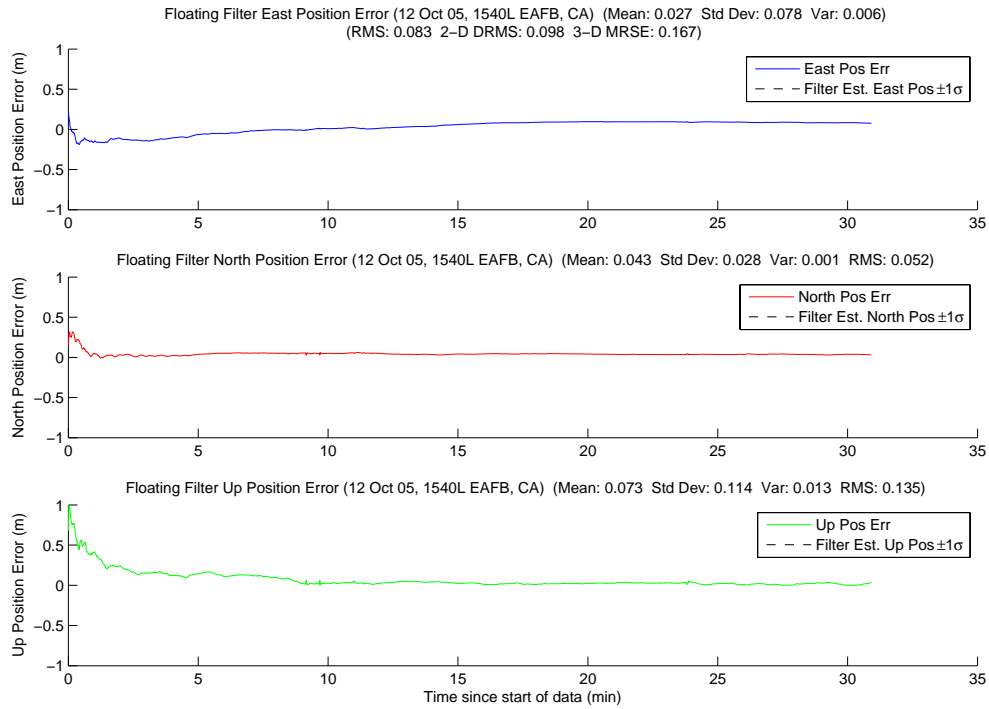


Figure A.301: Case F6.4: Floating Filter Position Errors

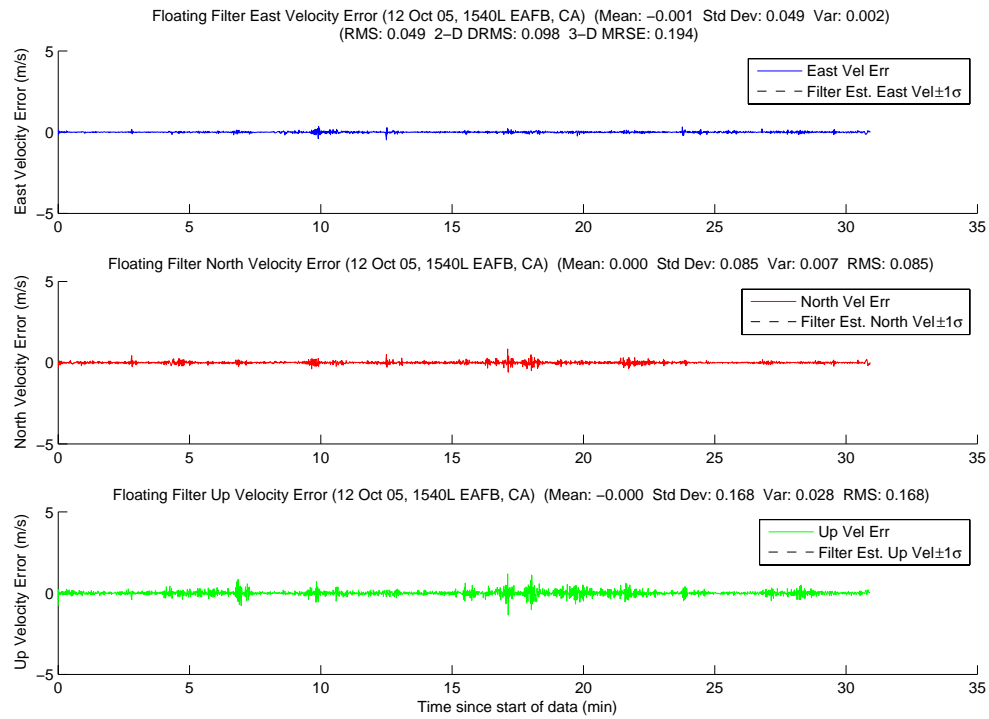


Figure A.302: Case F6.4: Floating Filter Velocity Errors

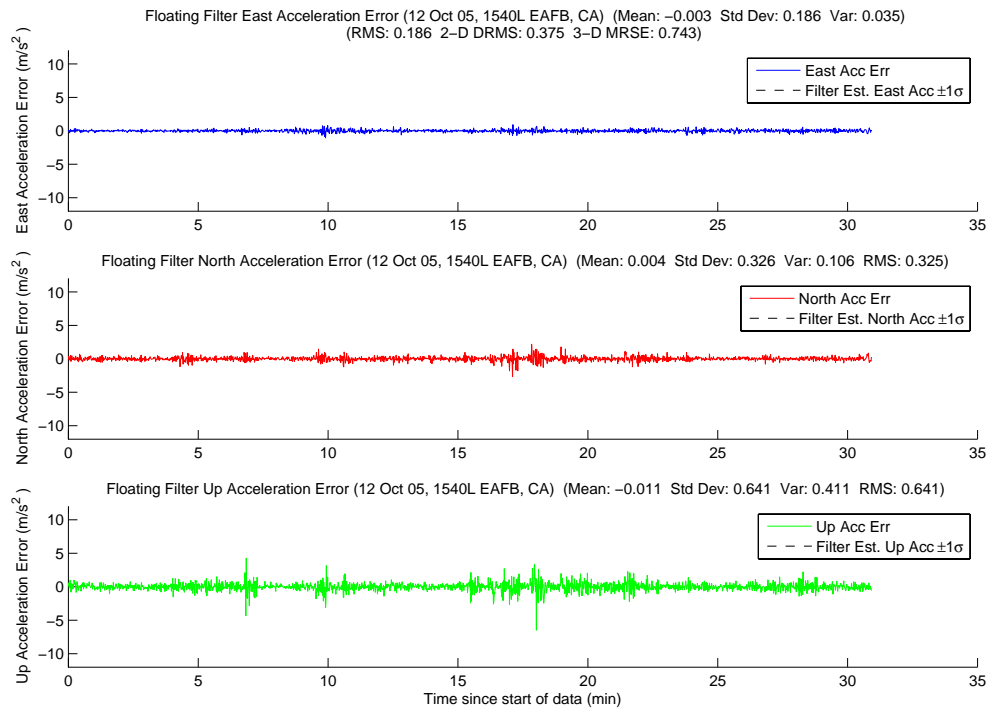


Figure A.303: Case F6.4: Floating Filter Acceleration Errors

Figure A.304 displays PRN 11 measurement information and the floating filter ambiguity estimate. All of the breaks in the floating filter's 1σ estimation were due to masking. PRN 11 was about to set, at 18.5 minutes, and was low on the horizon. The rest of the satellite's plots follow in Figures A.305, A.306, A.307, A.308, A.309, A.310, A.311, A.312, A.313, A.314, and A.315.

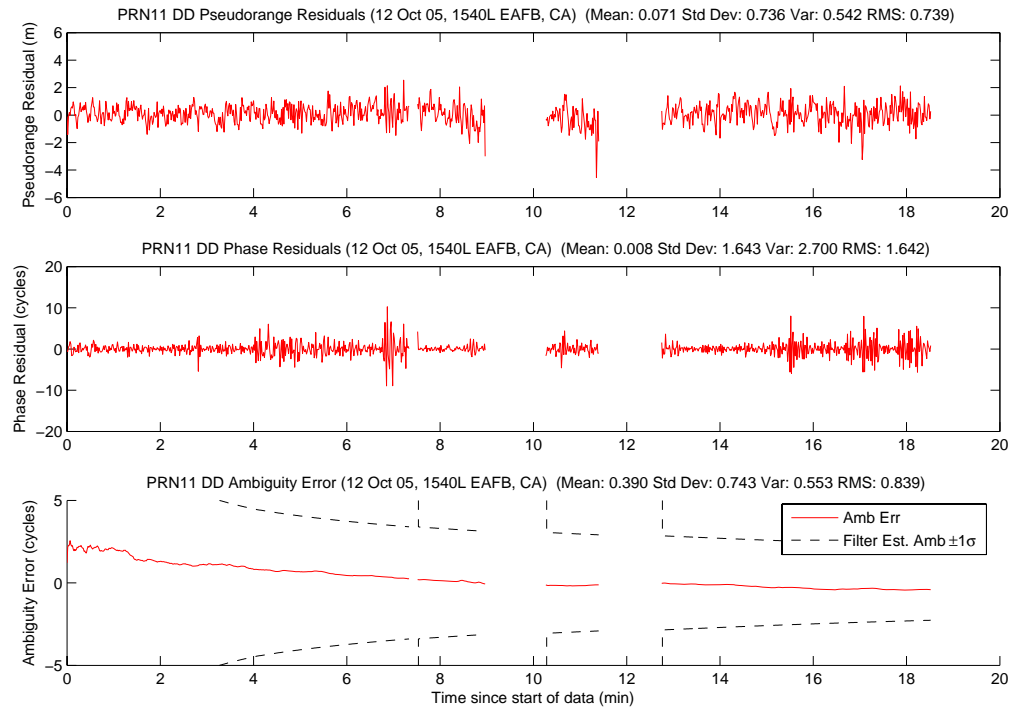


Figure A.304: Case F6.4: Satellite 11 Measurements

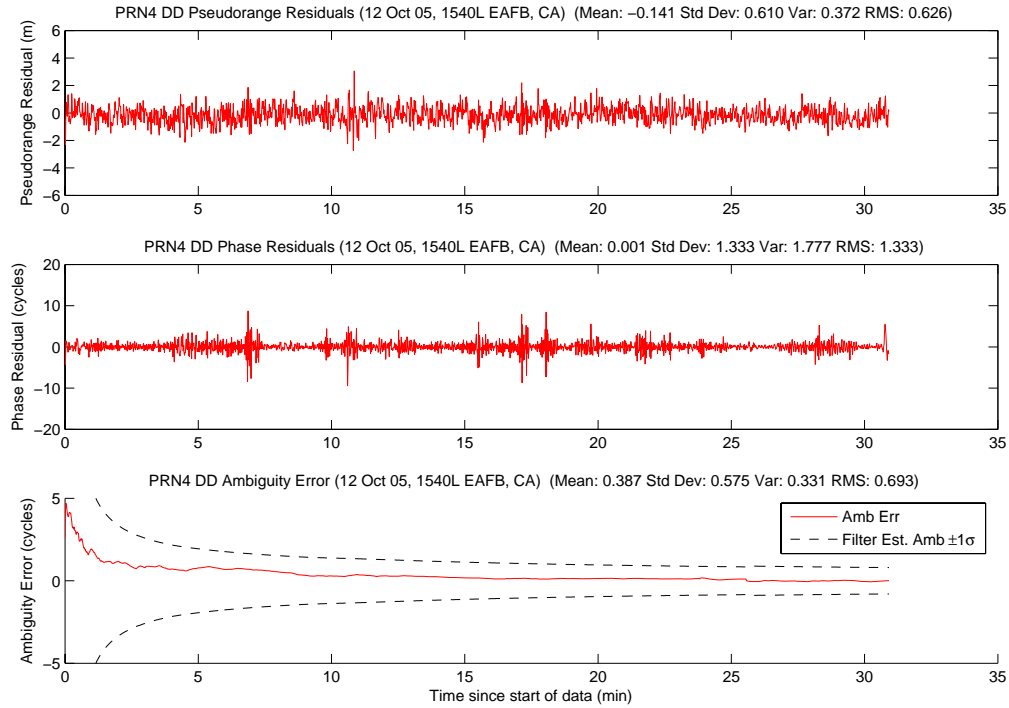


Figure A.305: Case F6.4: Satellite 4 Measurements

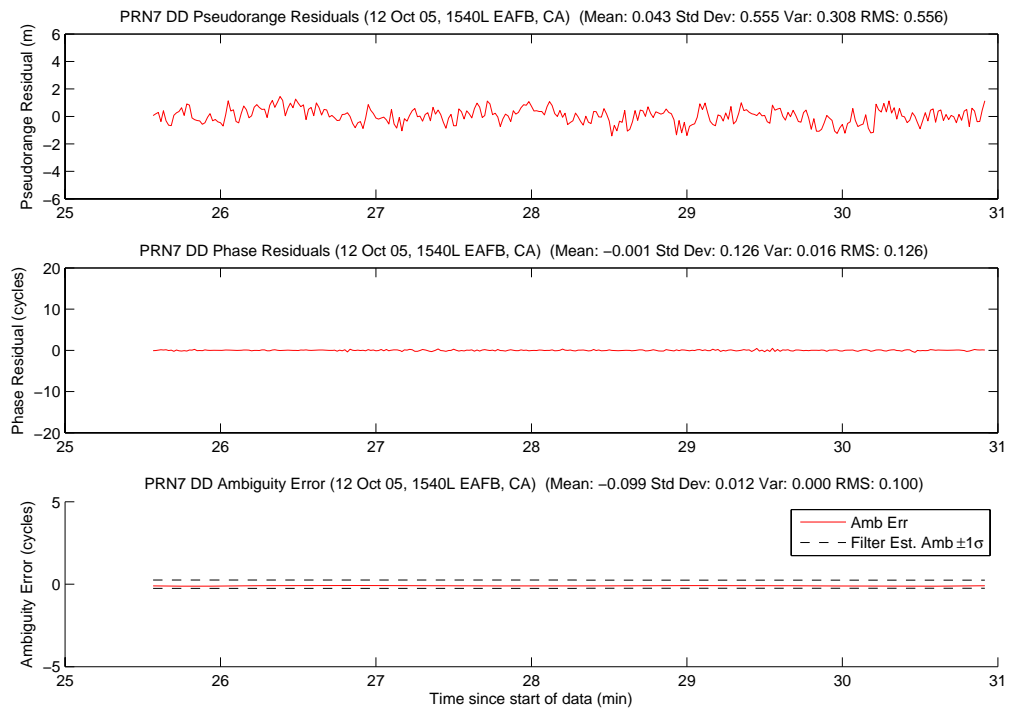


Figure A.306: Case F6.4: Satellite 7 Measurements

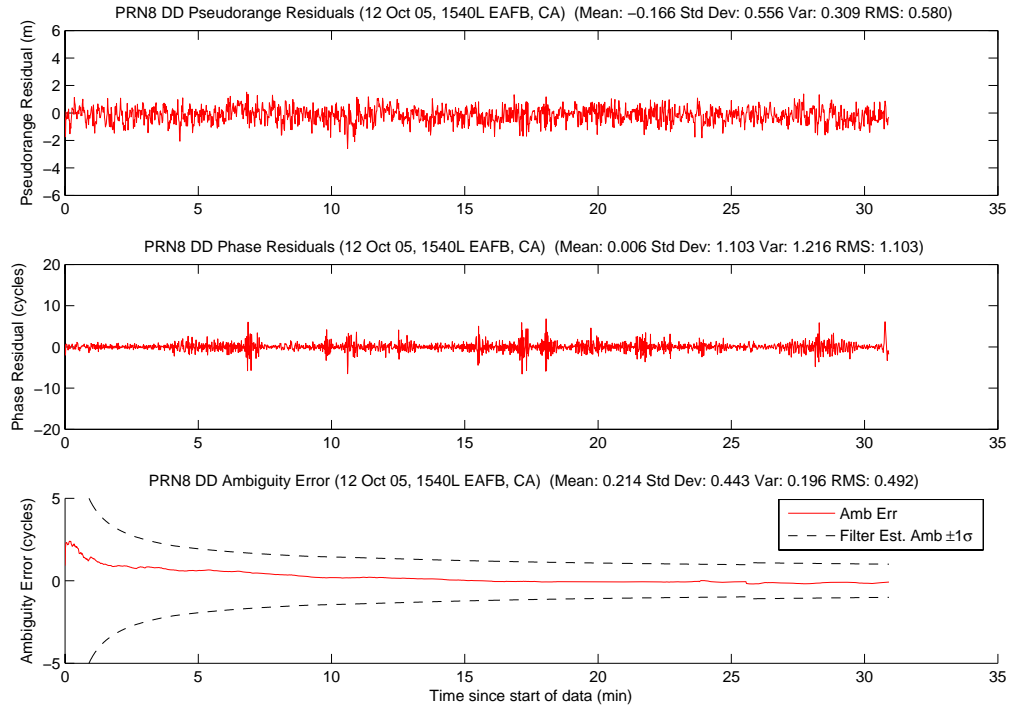


Figure A.307: Case F6.4: Satellite 8 Measurements

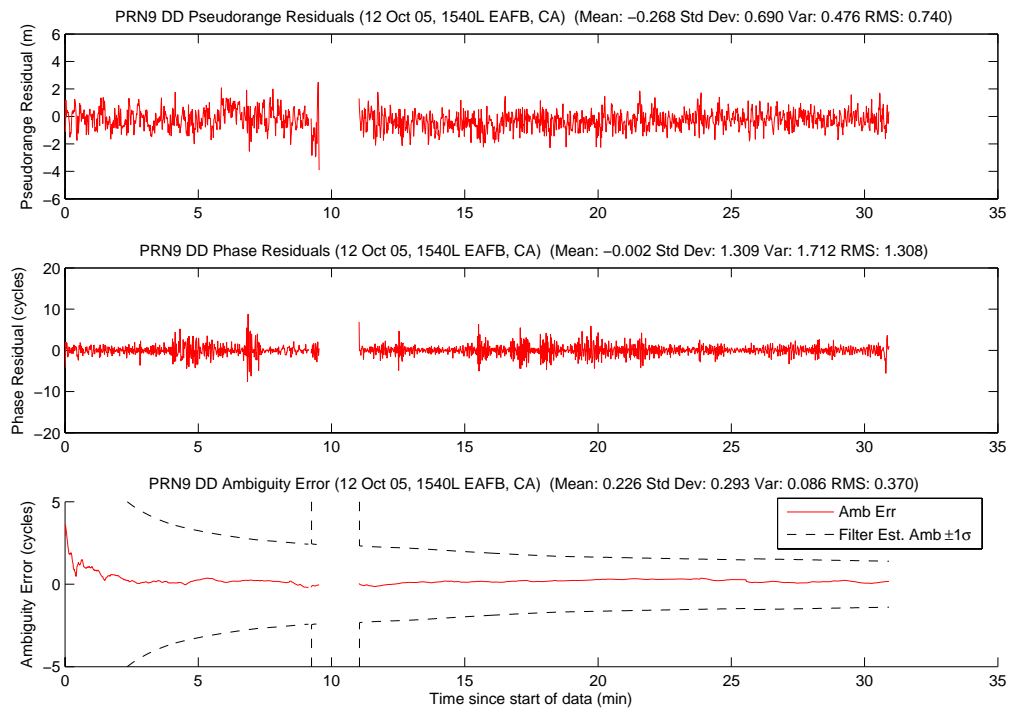


Figure A.308: Case F6.4: Satellite 9 Measurements

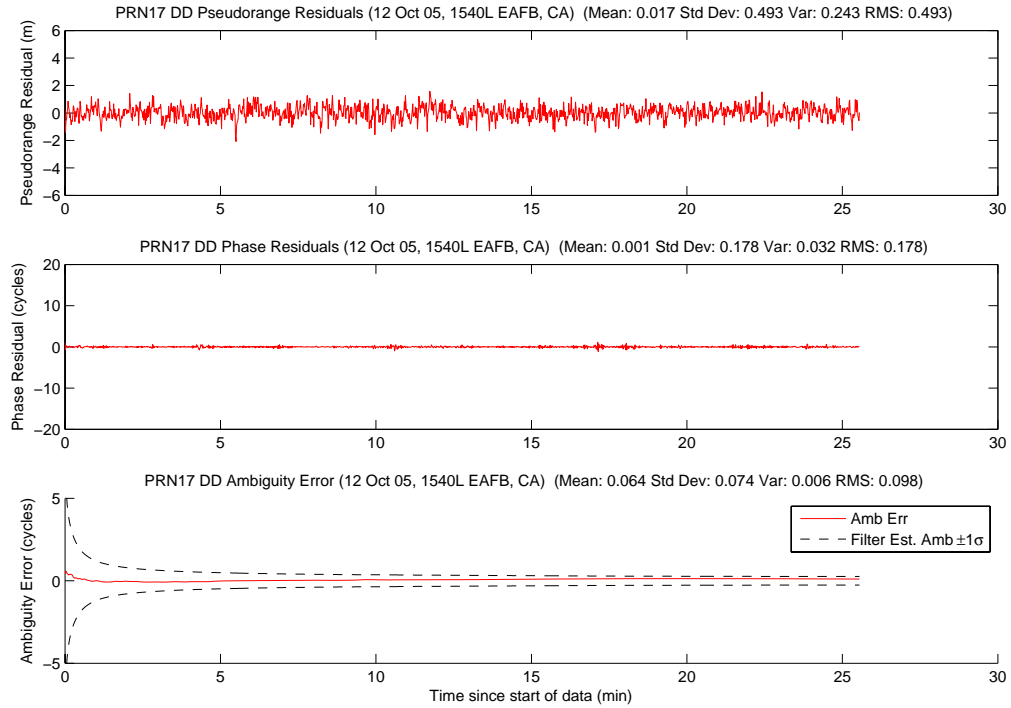


Figure A.309: Case F6.4: Satellite 17 Measurements

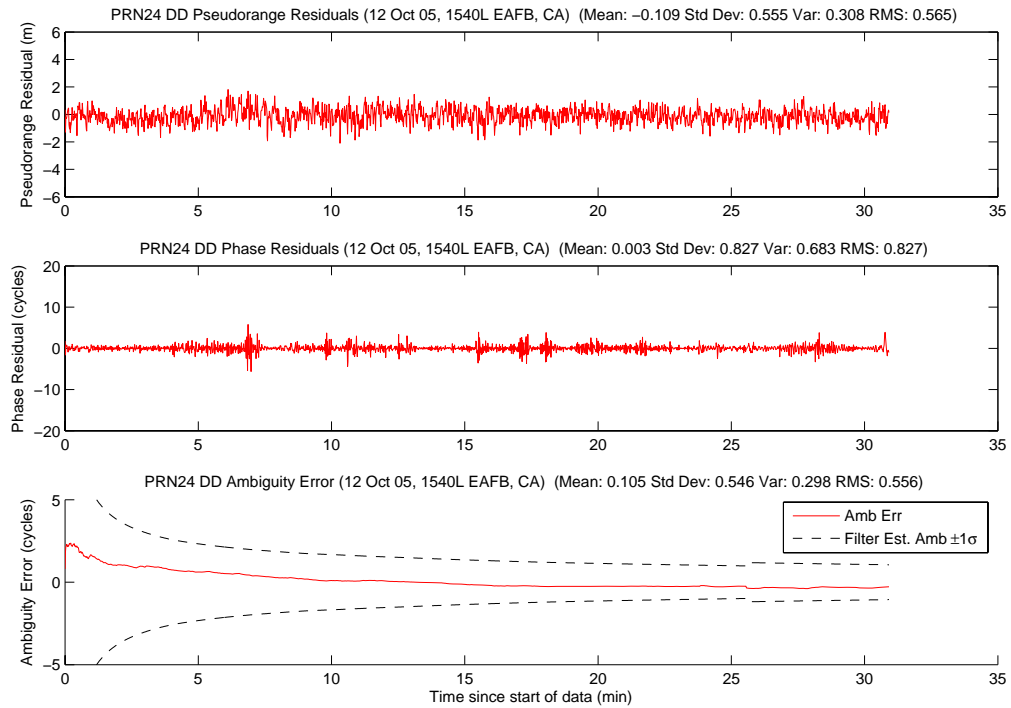


Figure A.310: Case F6.4: Satellite 24 Measurements

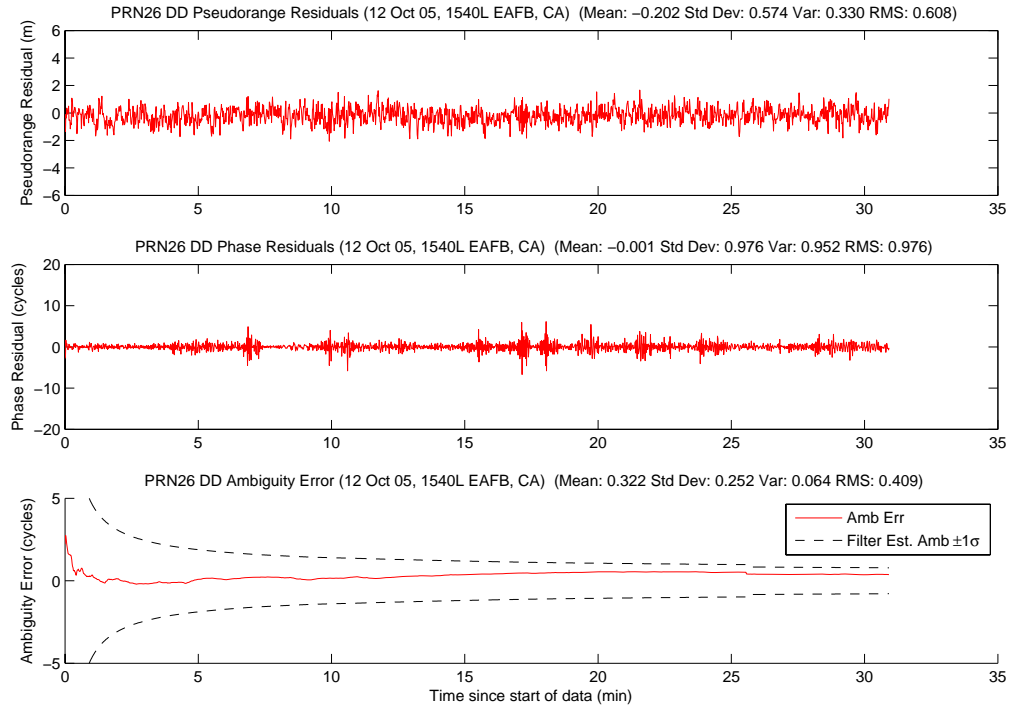


Figure A.311: Case F6.4: Satellite 26 Measurements

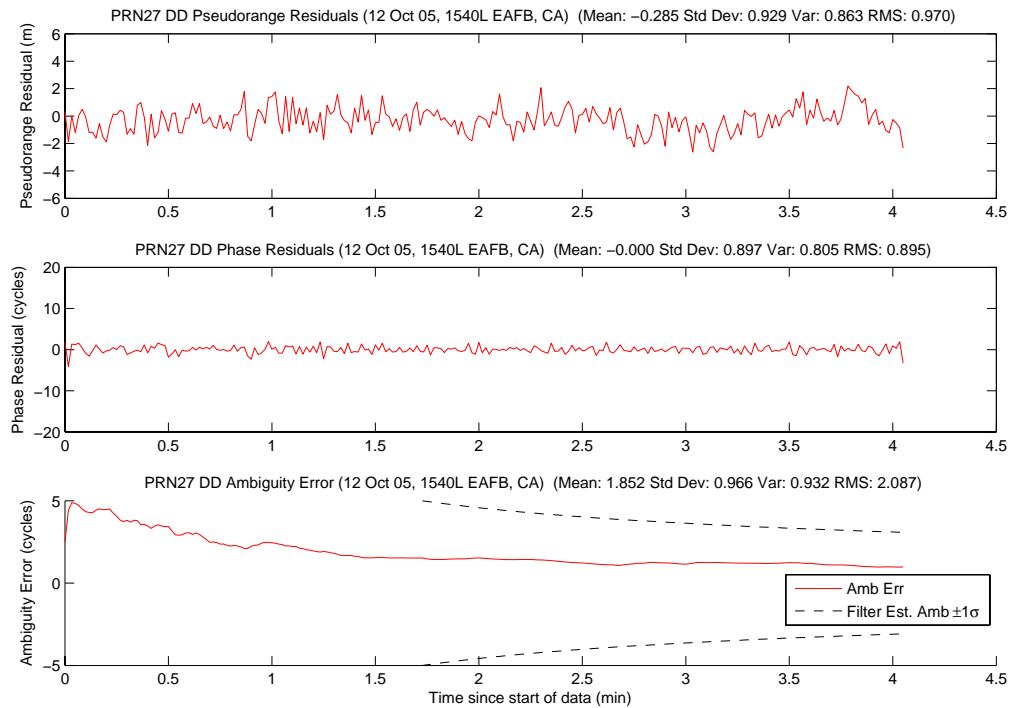


Figure A.312: Case F6.4: Satellite 27 Measurements

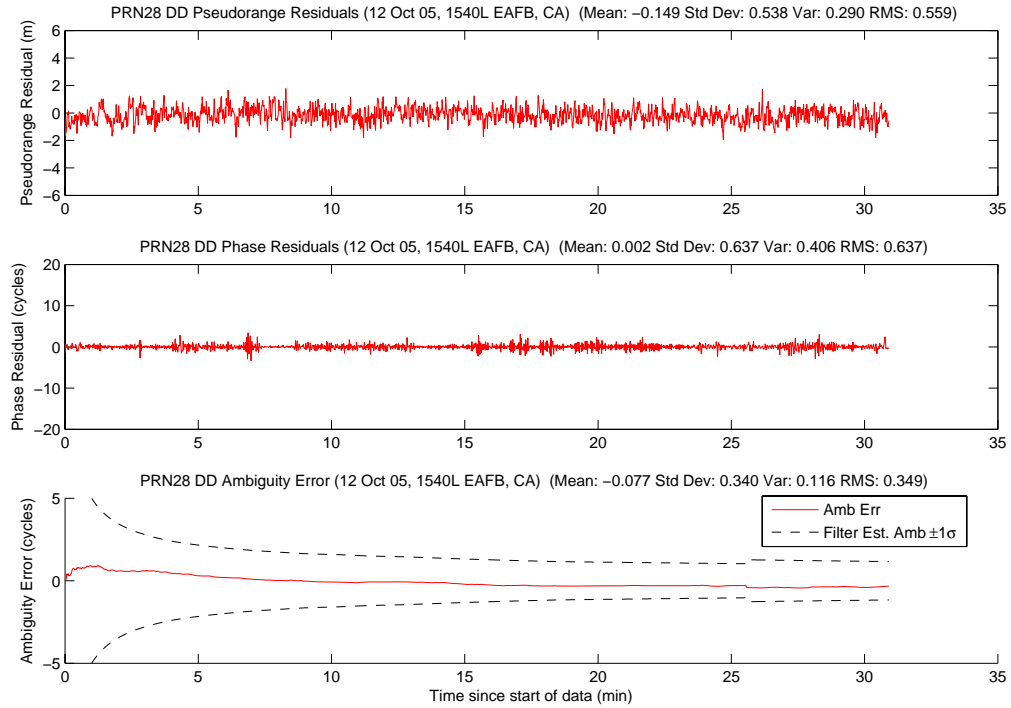


Figure A.313: Case F6.4: Satellite 28 Measurements

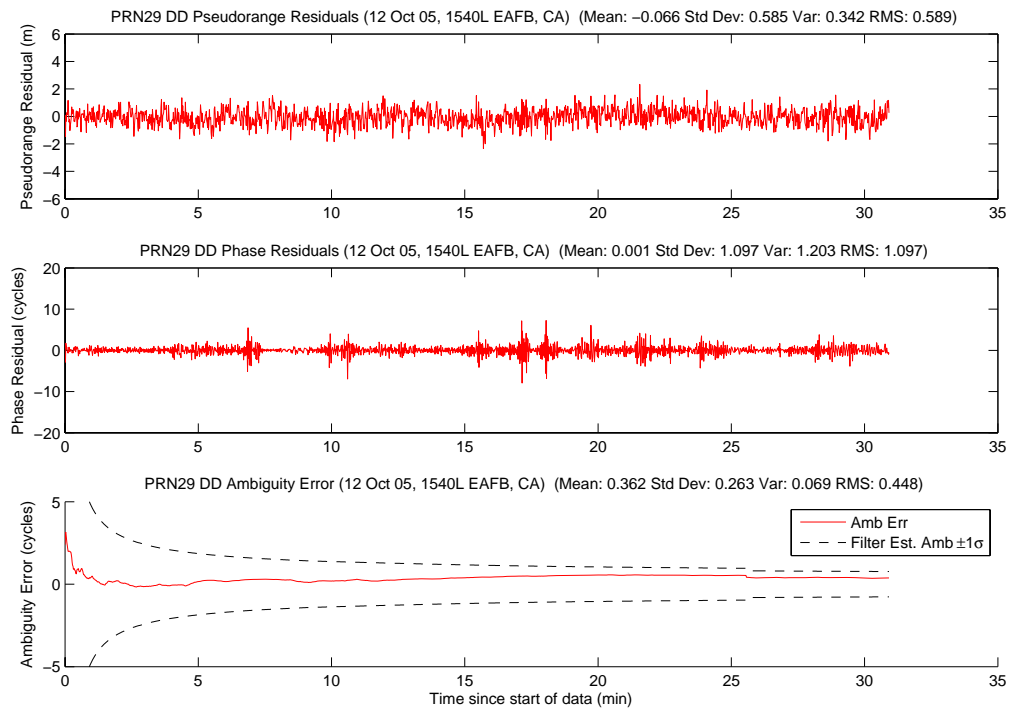


Figure A.314: Case F6.4: Satellite 29 Measurements

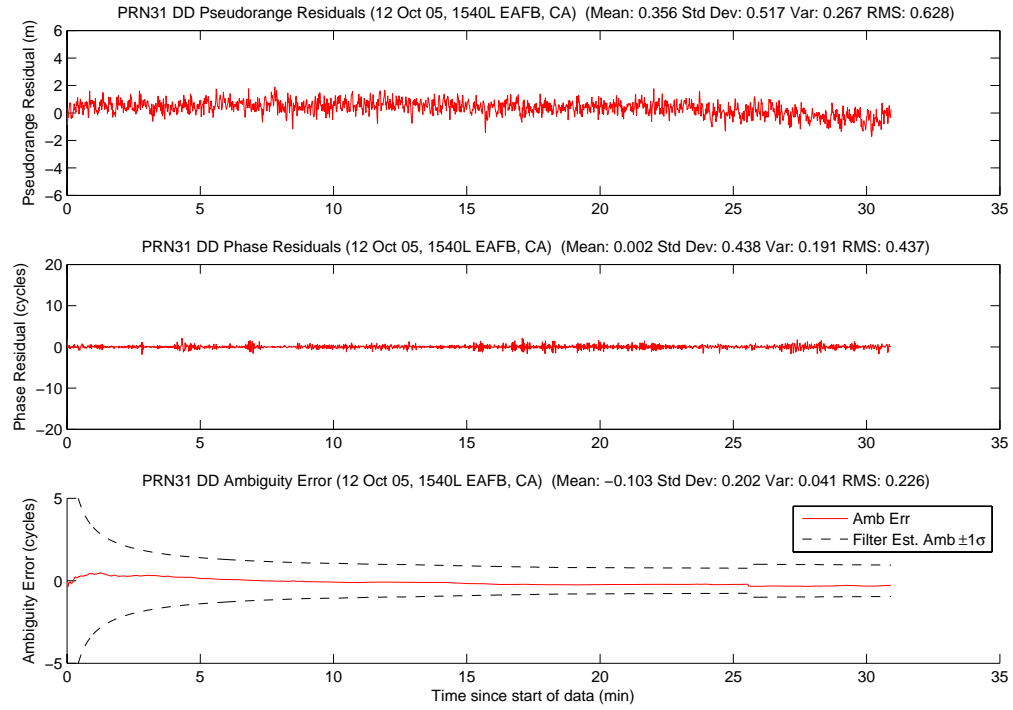


Figure A.315: Case F6.4: Satellite 31 Measurements

Figure A.316 shows the number of visible satellites. All of the satellite dropouts during the run were due to masking. Two satellites set during the run; PRN 27 at four minutes and PRN 11 at 18.5 minutes.

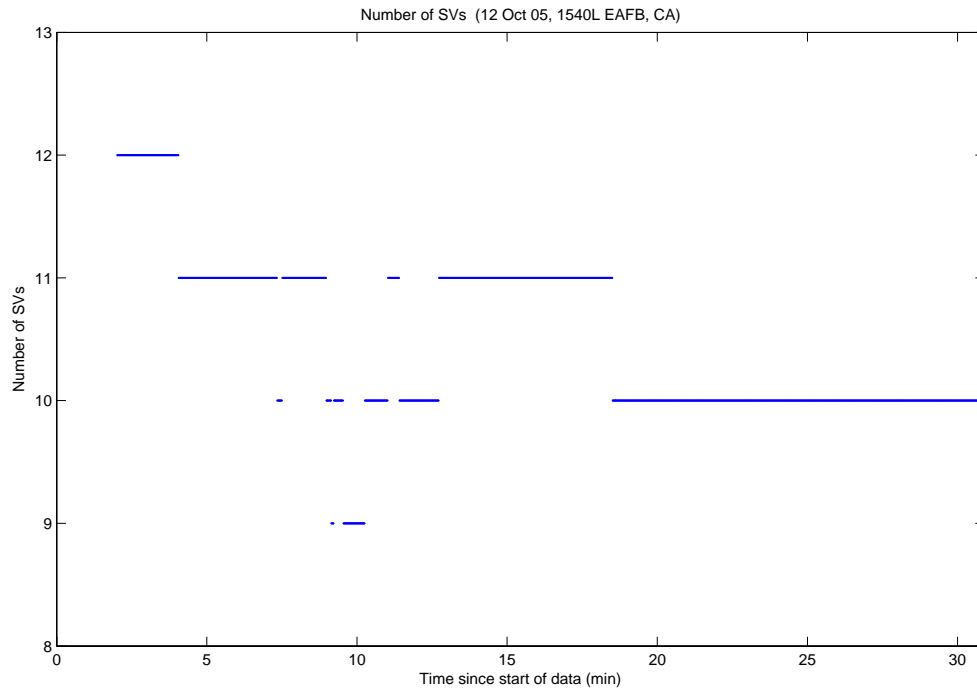


Figure A.316: Case F6.4: Visible Satellites

Figure A.317 shows all of the candidate ambiguity set's SOS residuals. The true ambiguity set's residuals are easily identifiable, even without the added emphasis. Figure A.318 displays the MMAE conditional probability. It is almost identical to Case 6.3.

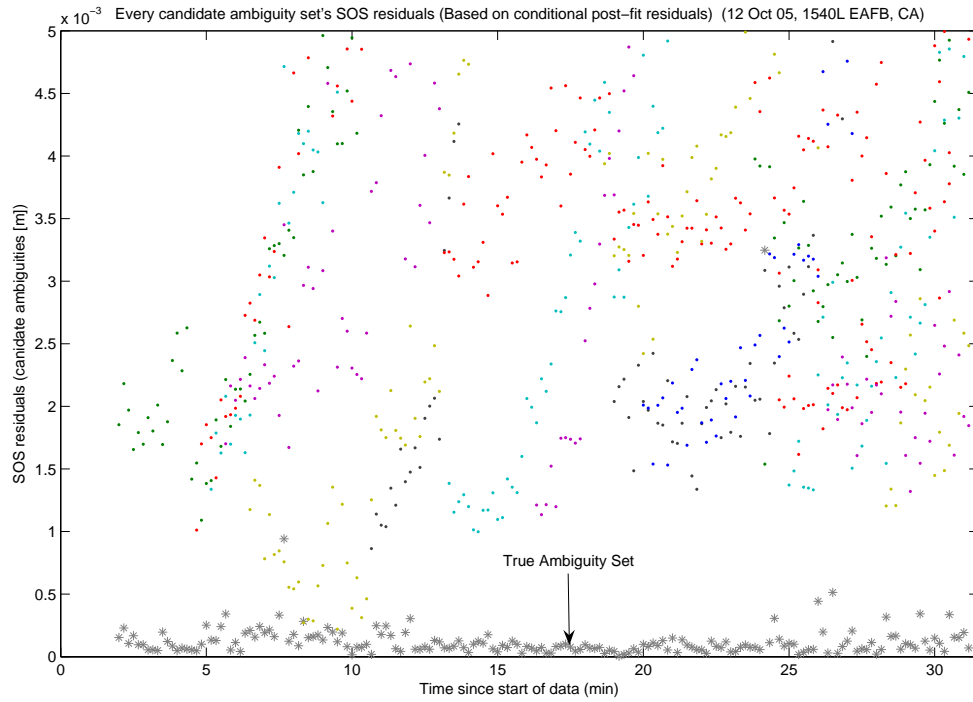


Figure A.317: Case F6.4: SOS Residuals

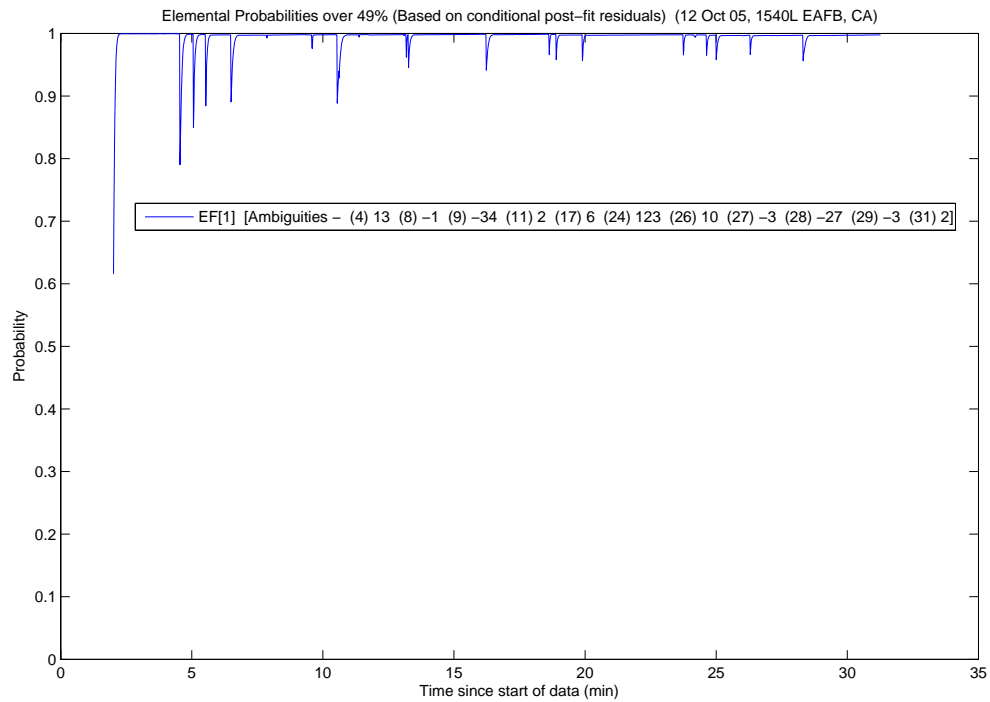


Figure A.318: Case F6.4: Conditional Probabilities for Selected MMAE Elemental Filters

The following figure shows the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.319 is EF[1].

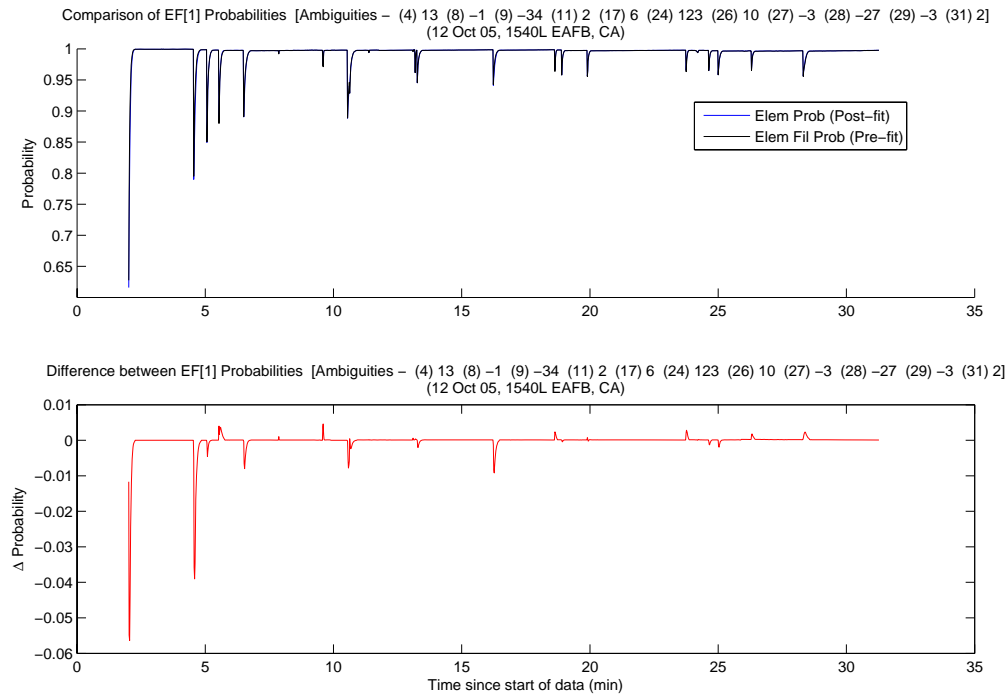


Figure A.319: Case F6.4: EF[1] Probability Comparison

The MMAE position error is shown in Figure A.320. The spikes in the position error correspond directly to the dips on the probability plot, Figure A.318.

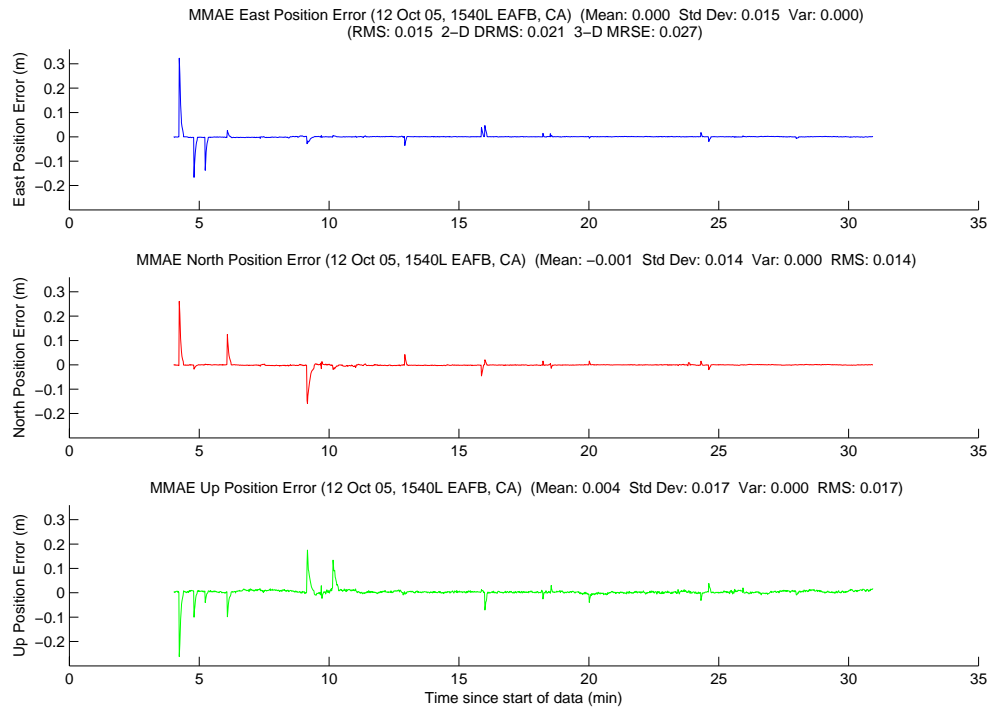


Figure A.320: Case F6.4: MMAE Position Error

Finally, Figure A.321 displays the position error plot for the minimum indicator variable. For the second time, the minimum indicator MRSE was within millimeters. Lastly, Table A.17 gives a summary of each method's position error.

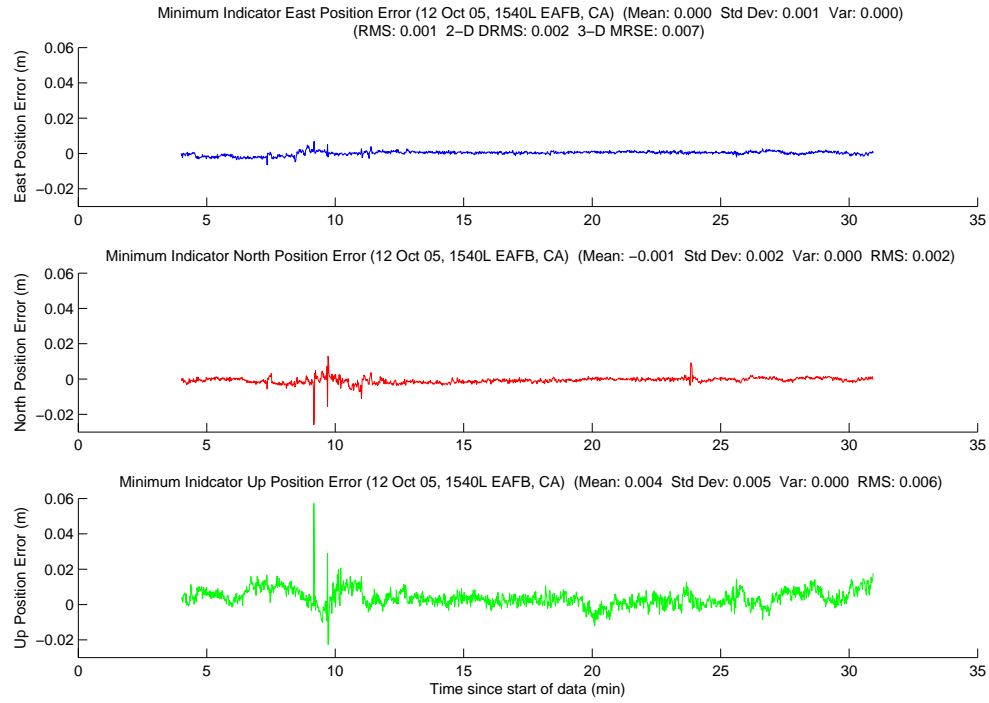


Figure A.321: Case F6.4: Minimum Indicator Position Error

Table A.17: Case F6.4: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.027	0.078	0.043	0.028	0.073	0.114	0.098	0.167
MMAE	0.000	0.015	-0.001	0.014	0.004	0.017	0.021	0.027
Min. Ind.	0.000	0.001	-0.001	0.002	0.003	0.005	0.002	0.007

A.2.14 Case F7.1: Flight Test 7, First Data Set . One test flight was flown on 13 October 2005. The first data run started at 1007L and lasted 59 minutes, 33 seconds. The weather was clear with light winds, and the aircraft ranged from 10.4 to 347.4 meters apart during the run. Figure A.322 shows the minimum indicator North, East, Down relative positions.

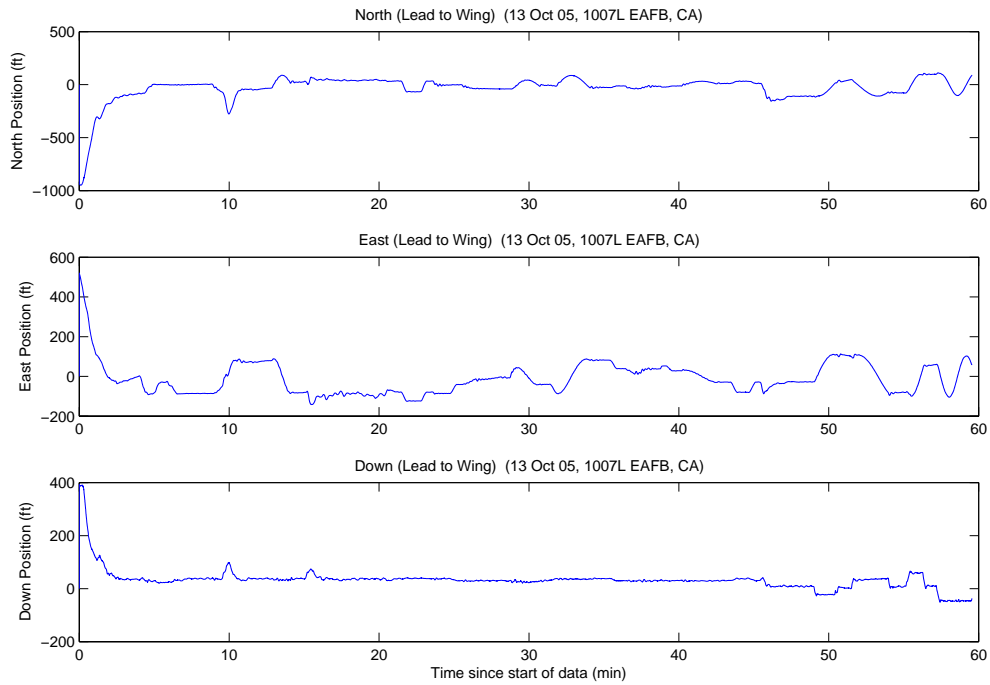


Figure A.322: Case F7.1: North, East, Down Relative Positions

A total of 25 unique ambiguity sets were returned from the LAMBDA function during the data run. The aircraft were 347.4 meters apart when the system started, with eight visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 3453 of 3453 epochs (100.0 percent), with a maximum ratio of 142.8.

Figures A.323, A.324, and A.325 show the floating filter errors. The floating filter performed as expected. The floating filter's 1σ estimate of the position error is outside the plot scales in all three figures.

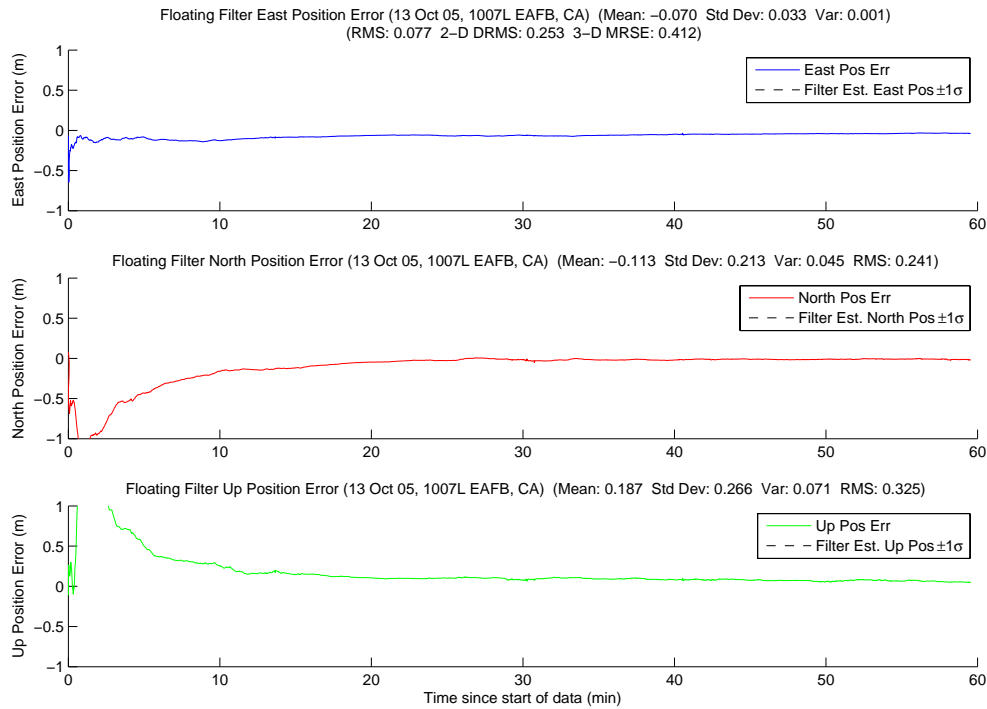


Figure A.323: Case F7.1: Floating Filter Position Errors

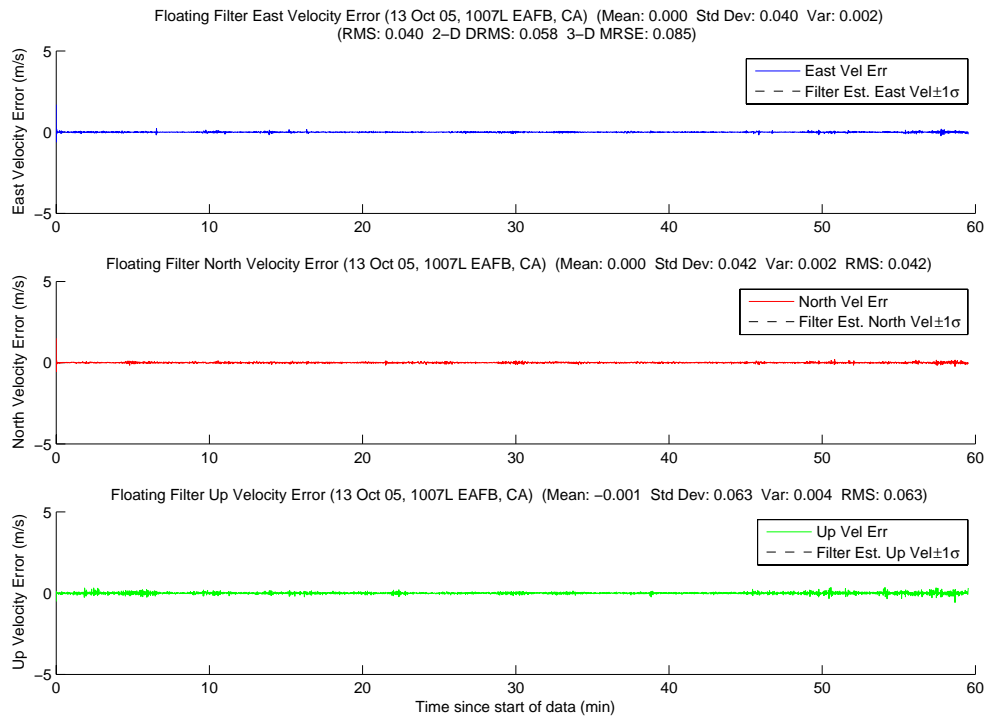


Figure A.324: Case F7.1: Floating Filter Velocity Errors

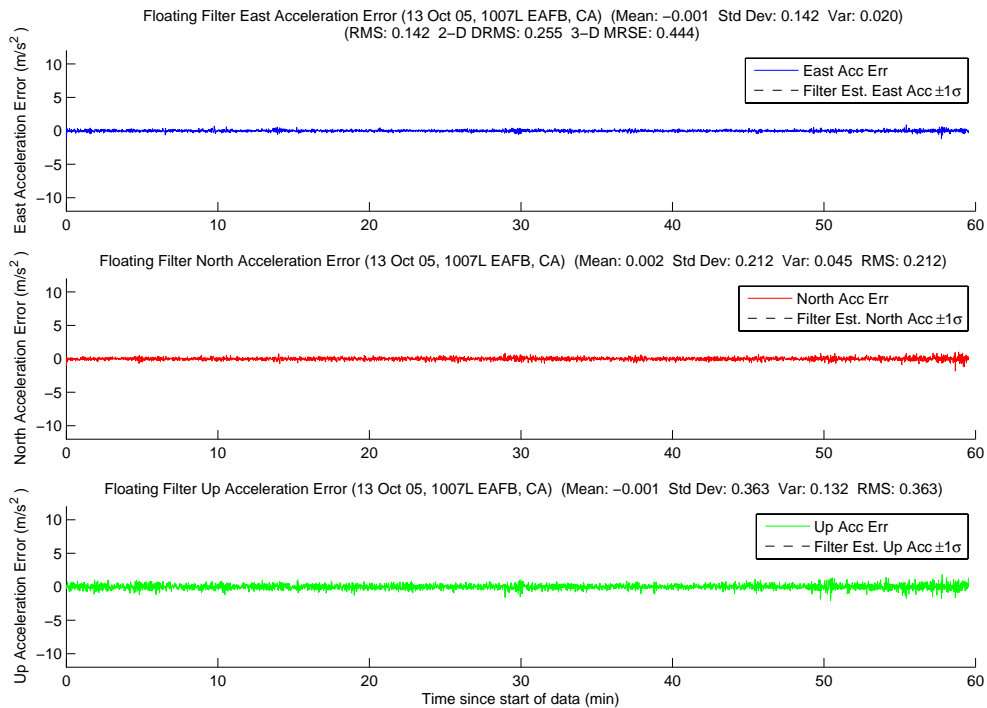


Figure A.325: Case F7.1: Floating Filter Acceleration Errors

Figure A.326 displays PRN 16 measurement information and the floating filter ambiguity estimate. The break in the floating filter's 1σ estimation at 29 minutes was due to masking. The rest of the satellite's plots follow in Figures A.327, A.328, A.329, A.330, A.331, A.332, A.333, and A.334.

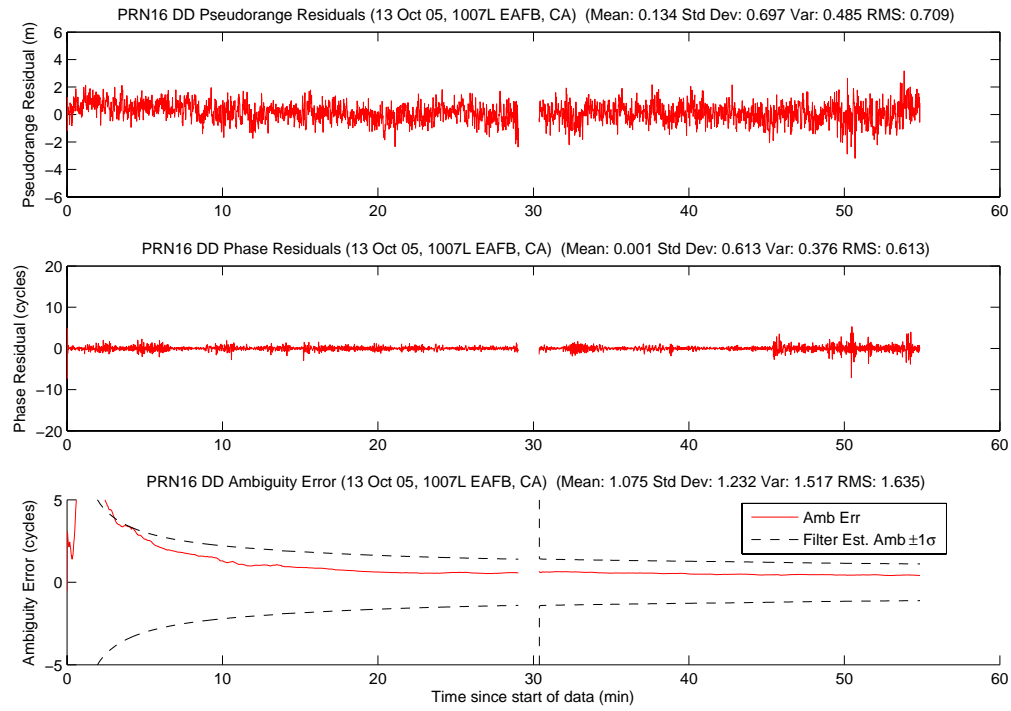


Figure A.326: Case F7.1: Satellite 16 Measurements

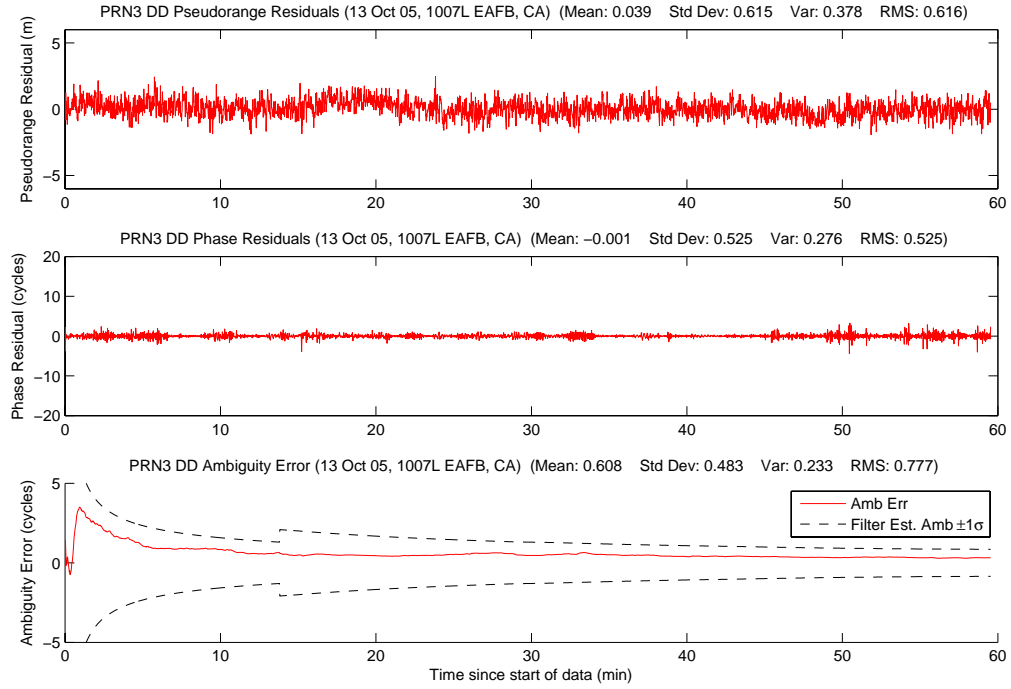


Figure A.327: Case F7.1: Satellite 3 Measurements

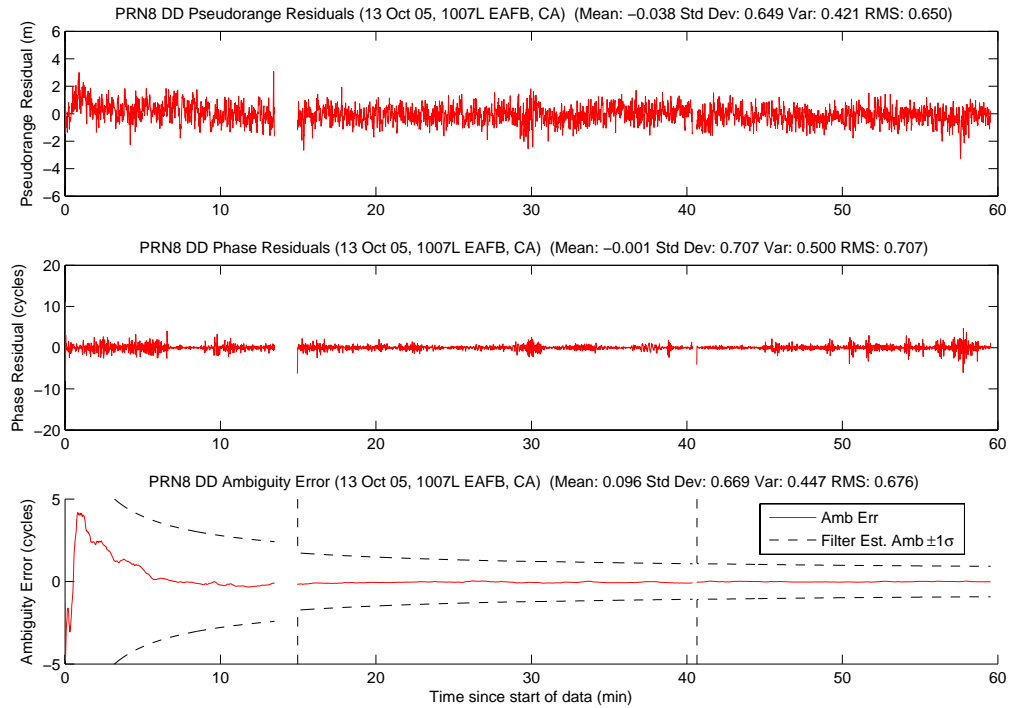


Figure A.328: Case F7.1: Satellite 8 Measurements

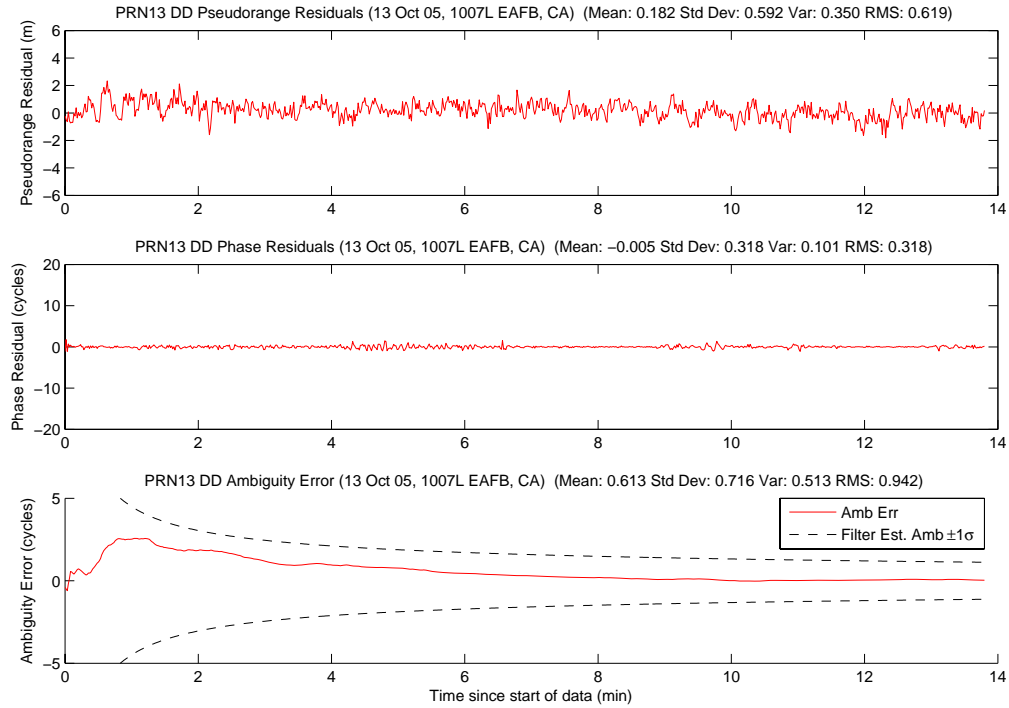


Figure A.329: Case F7.1: Satellite 13 Measurements

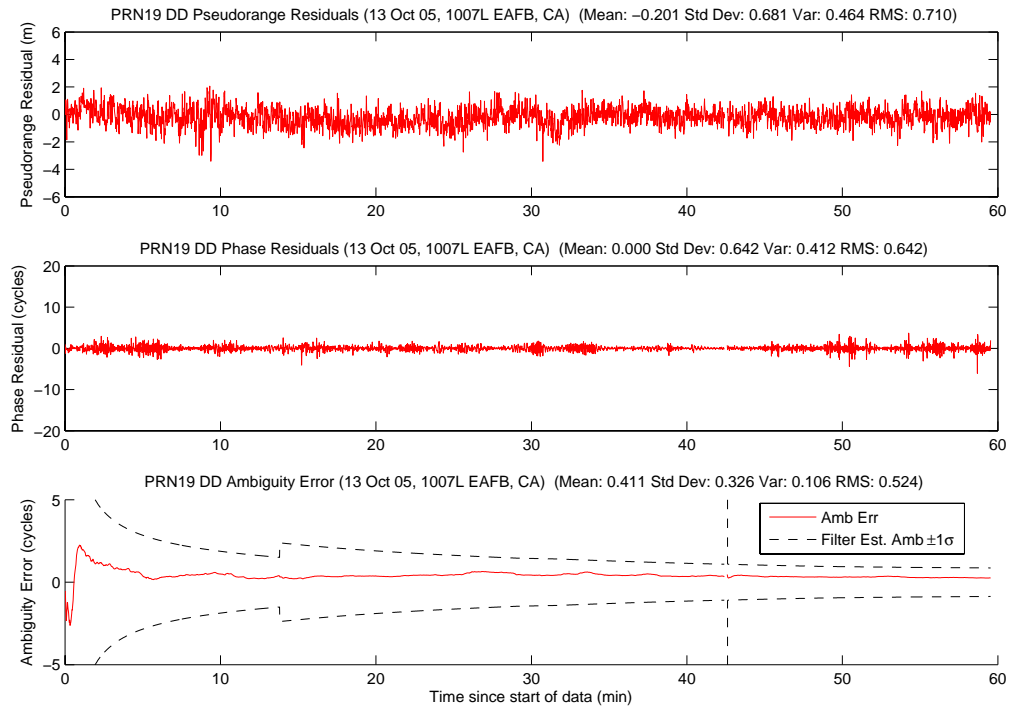


Figure A.330: Case F7.1: Satellite 19 Measurements

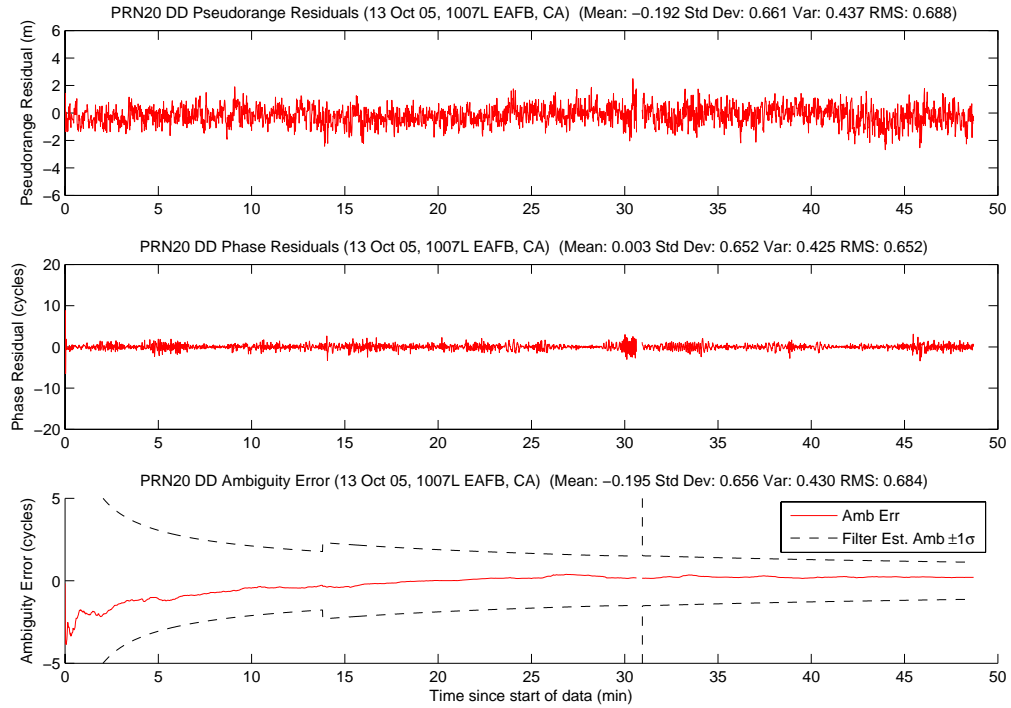


Figure A.331: Case F7.1: Satellite 20 Measurements

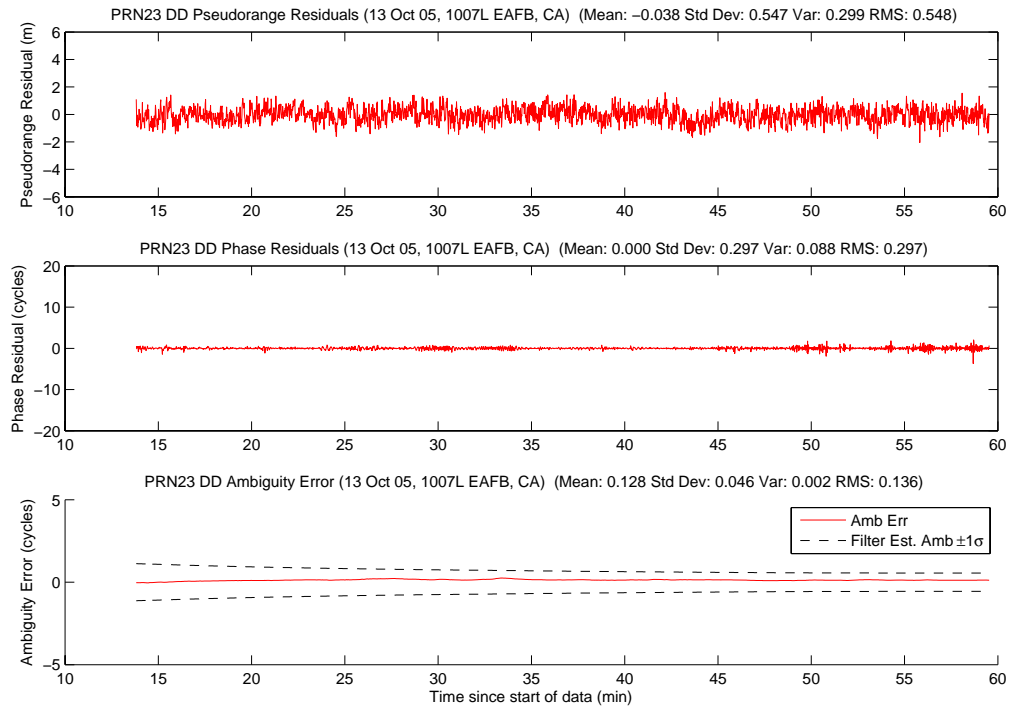


Figure A.332: Case F7.1: Satellite 23 Measurements

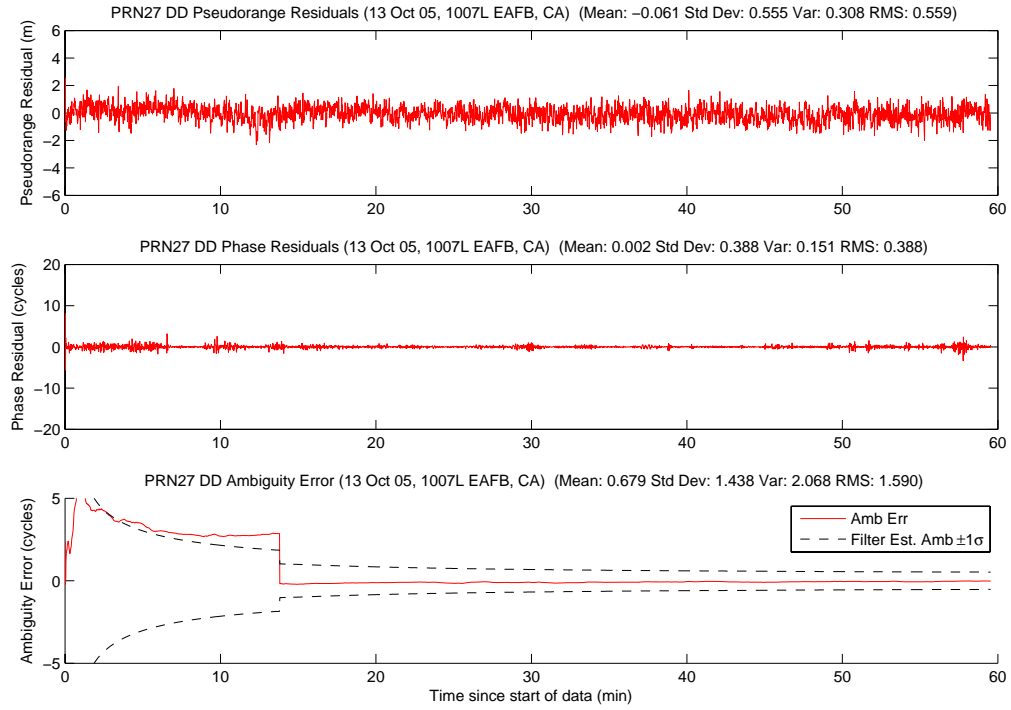


Figure A.333: Case F7.1: Satellite 27 Measurements

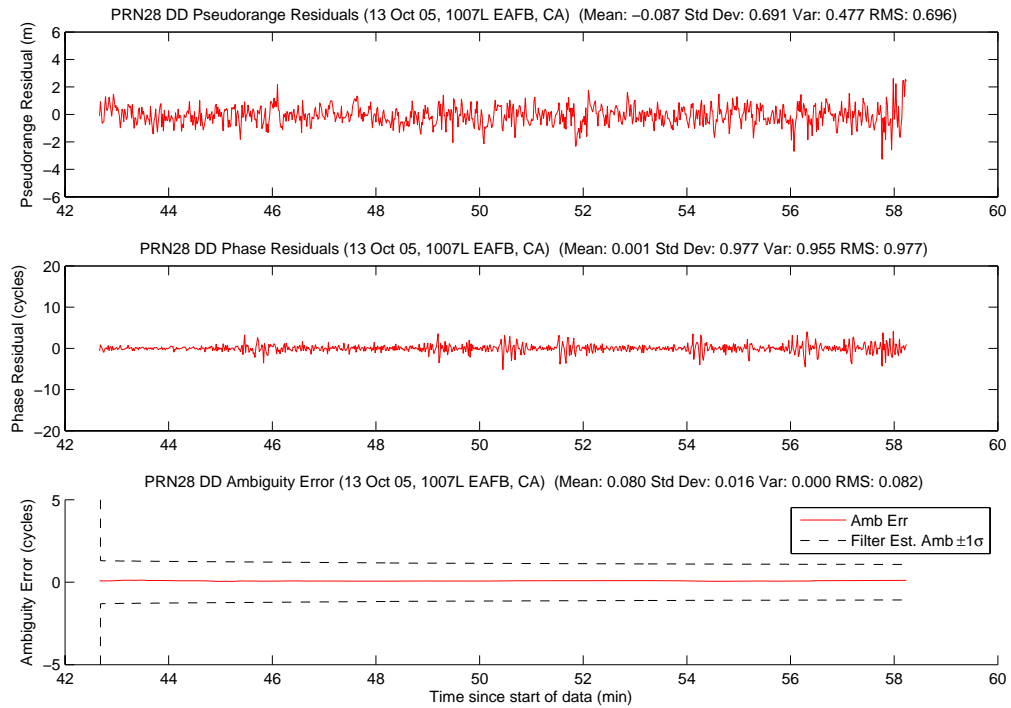


Figure A.334: Case F7.1: Satellite 28 Measurements

Figure A.335 shows the number of visible satellites. There were multiple satellite dropouts during the run, and all were due to satellite masking. PRN 8 masked at 13.5 and 40.5 minutes, PRN 16 at 29 minutes, PRN 20 at 30.5 minutes, PRN 19 at 42.5 minutes, and PRN 28 at 58 minutes. One satellite rose, PRN 28, at 42.5 minutes, and two satellites set, PRN 20 and 16, at 48.7 and 55 minutes, respectively.

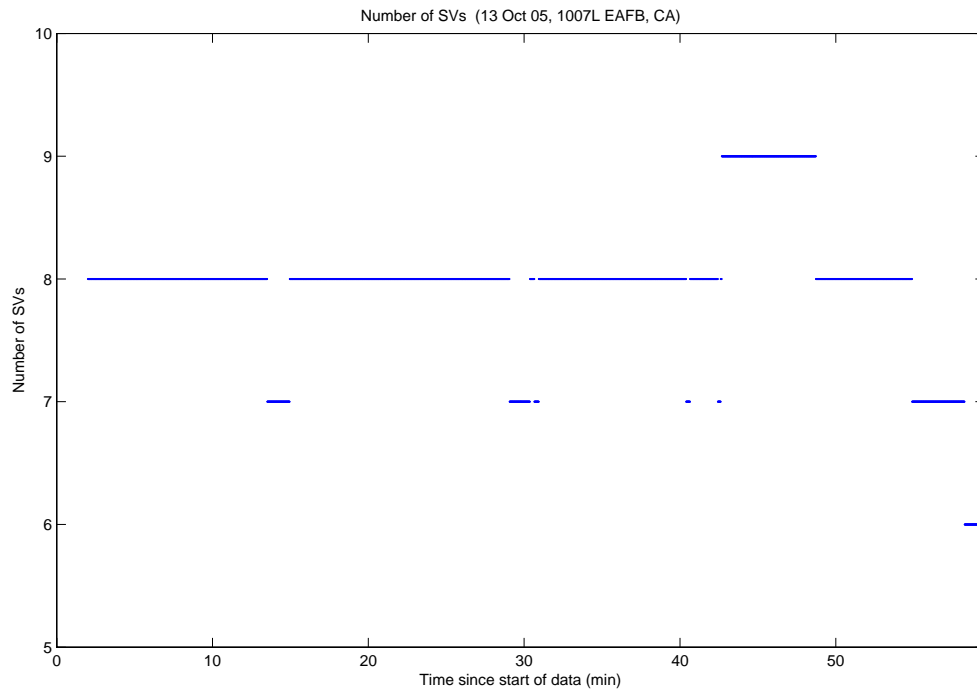


Figure A.335: Case F7.1: Visible Satellites

Figure A.336 shows all of the candidate ambiguity set's SOS residuals. The residual plot is similar to the other plots with eight visible satellites. Figure A.337 displays the MMAE conditional probability. Interestingly, the true set's probability did not dip, even at the end of the run, when there were only six visible satellites. The MMAE had garnered so much information on the system dynamics from previous measurements, it was not “confused” by erroneous ambiguity sets. There were a couple of erroneous data set's residuals at or below the true set's residuals in Figure A.336 at 58 minutes.

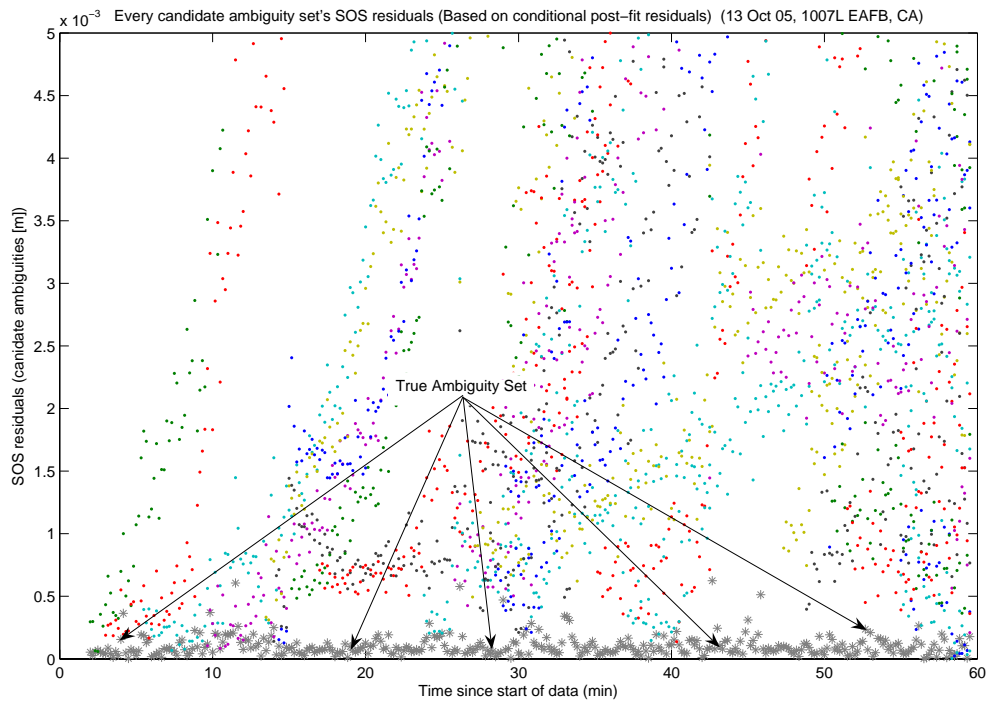


Figure A.336: Case F7.1: SOS Residuals

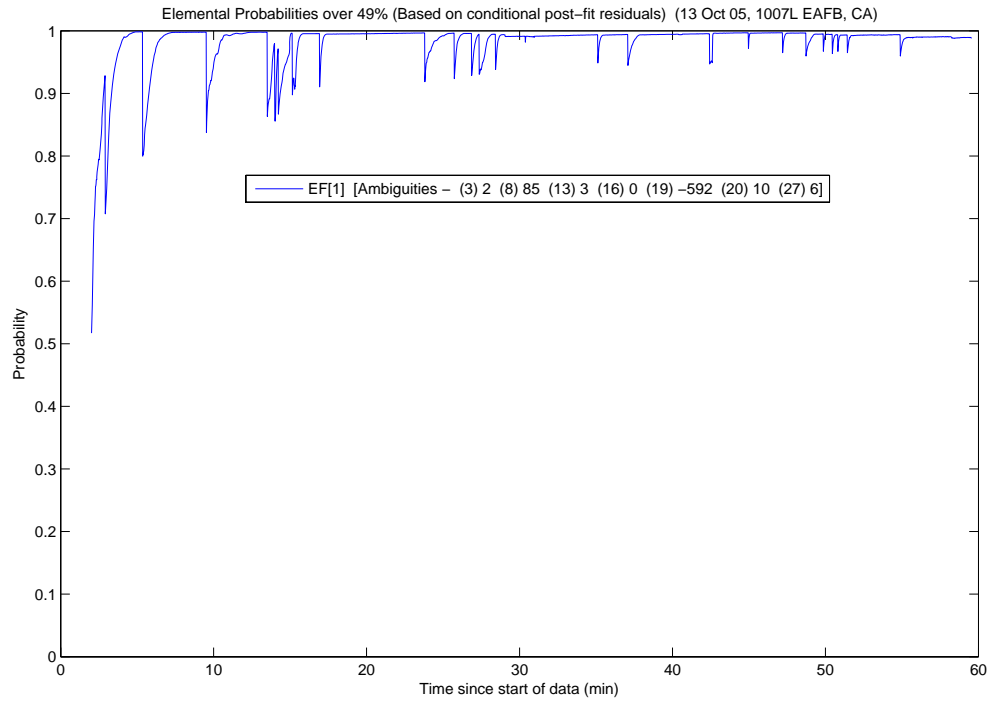


Figure A.337: Case F7.1: Conditional Probabilities for Selected MMAE Elemental Filters

The following figure shows the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.338 is EF[1].

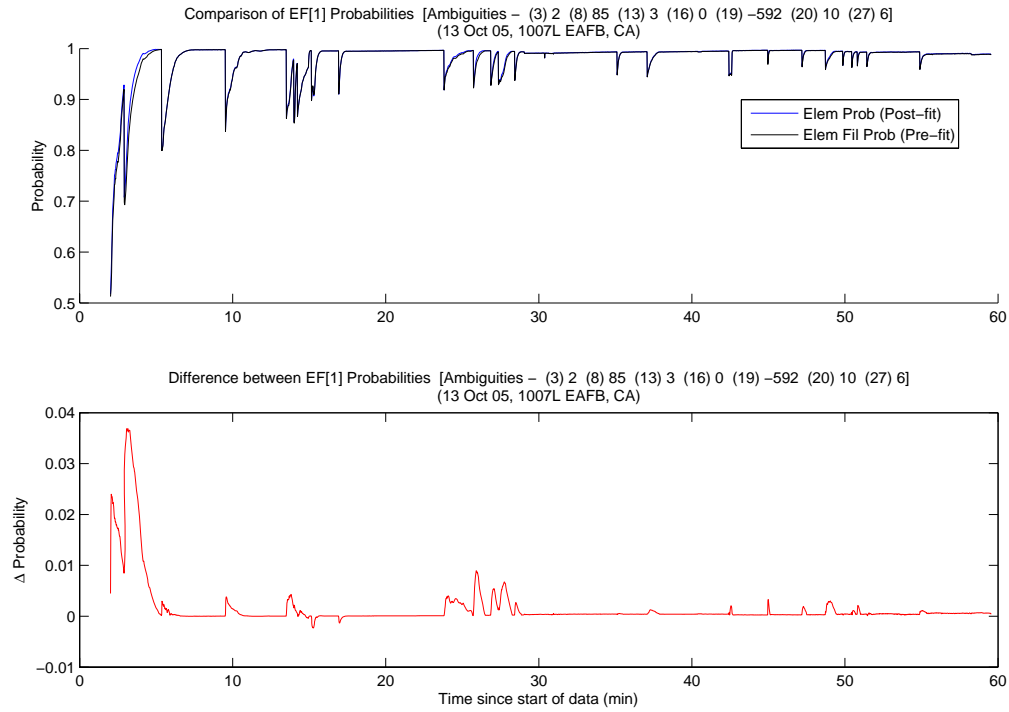


Figure A.338: Case F7.1: EF[1] Probability Comparison

The MMAE position error is shown in Figure A.339. The MMAE performed extremely well. The few spikes in the position error corresponds to places on the probability plot where the true ambiguity set's probability dipped.

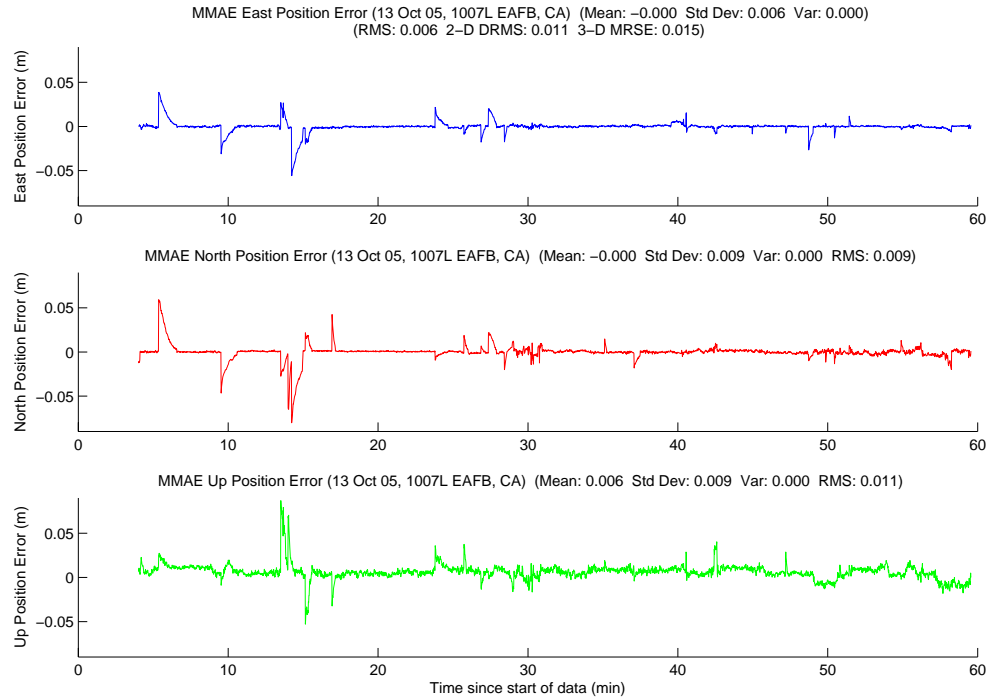


Figure A.339: Case F7.1: MMAE Position Error

Finally, Figure A.340 displays the position error plot for the minimum indicator variable. The minimum indicator position solution was within millimeters. It appears the incoming phase measurements were noisy around 30 minutes. This also corresponded to two different satellites masking. The aircraft were most like maneuvering and there was possibly some very small phase multipath between the aircraft. The other spikes in the position error at 15 and 40 minutes also corresponded to satellite masking. Lastly, Table A.18 gives a summary of each method's position error.

Table A.18: Case F7.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.070	0.033	-0.113	0.213	0.187	0.266	0.253	0.412
MMAE	-0.000	0.006	-0.000	0.009	0.006	0.009	0.011	0.015
Min. Ind.	-0.000	0.001	0.000	0.002	0.006	0.006	0.002	0.008

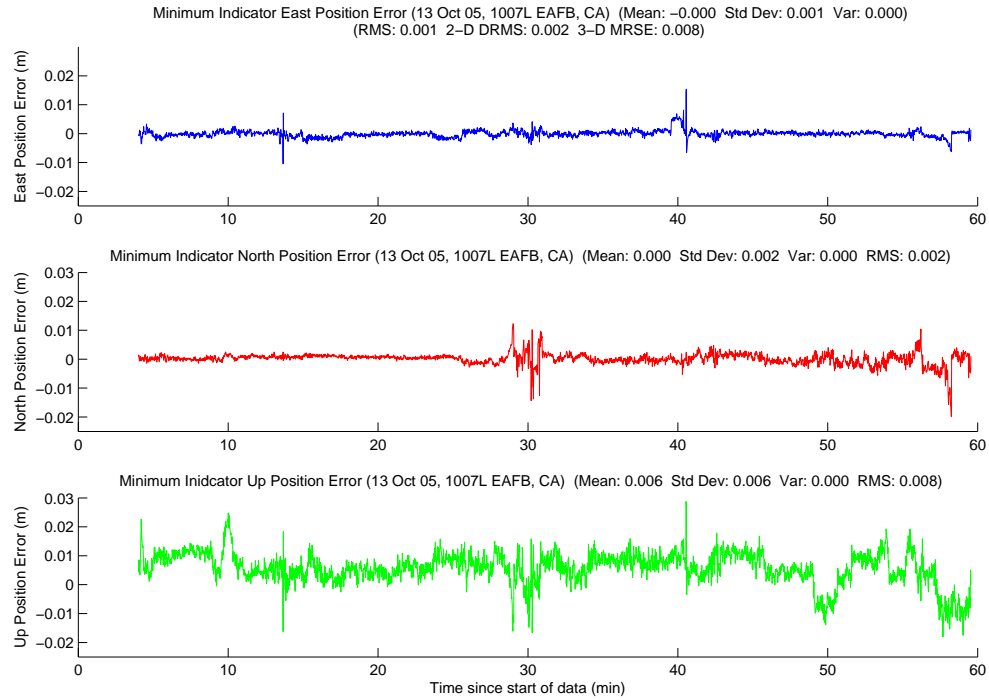


Figure A.340: Case F7.1: Minimum Indicator Position Error

A.2.15 Case F7.2: Flight Test 7, Second Data Set . Case 7.2 started at 1107L, immediately after the end of Case 7.1. The run lasted 41 minutes, 34 seconds and the aircraft ranged from 11.6 to 342.1 meters apart during the data run. Figure A.341 shows the minimum indicator North, East, Down relative positions.

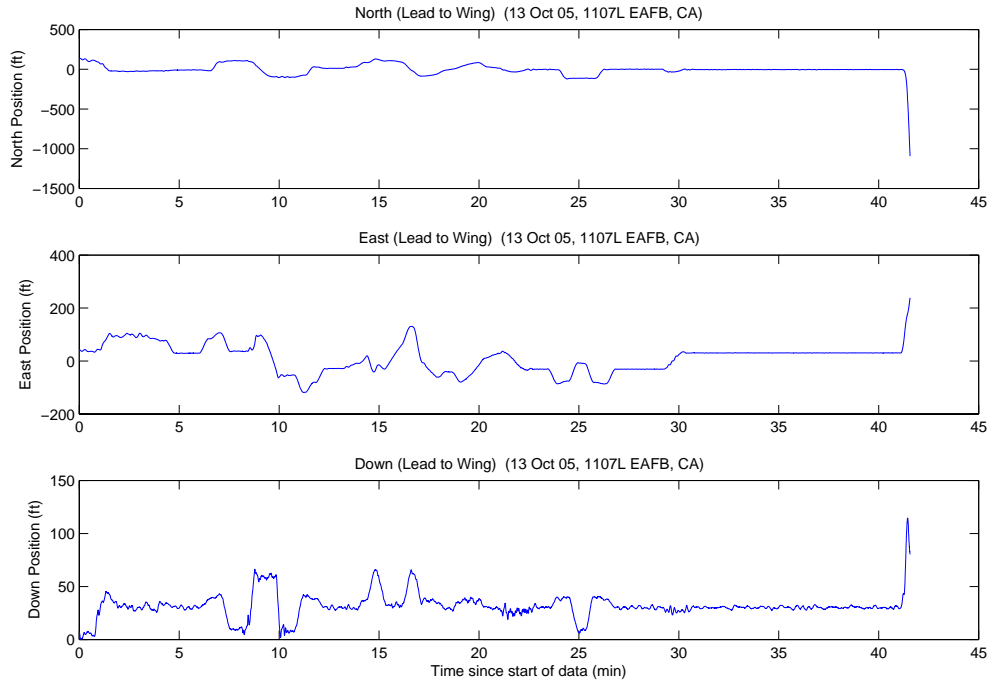


Figure A.341: Case F7.2: North, East, Down Relative Positions

A total of 25 unique ambiguity sets were returned from the LAMBDA function during the data run. The aircraft were 42.8 meters apart when the system started, with seven visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 2395 of 2395 epochs (100.0 percent), with a maximum ratio of 49.3.

Figures A.342, A.343, and A.344 show the floating filter position errors. The floating filter performed as expected. The floating filter's 1σ estimate of the position error is outside the plot scales in all three figures.

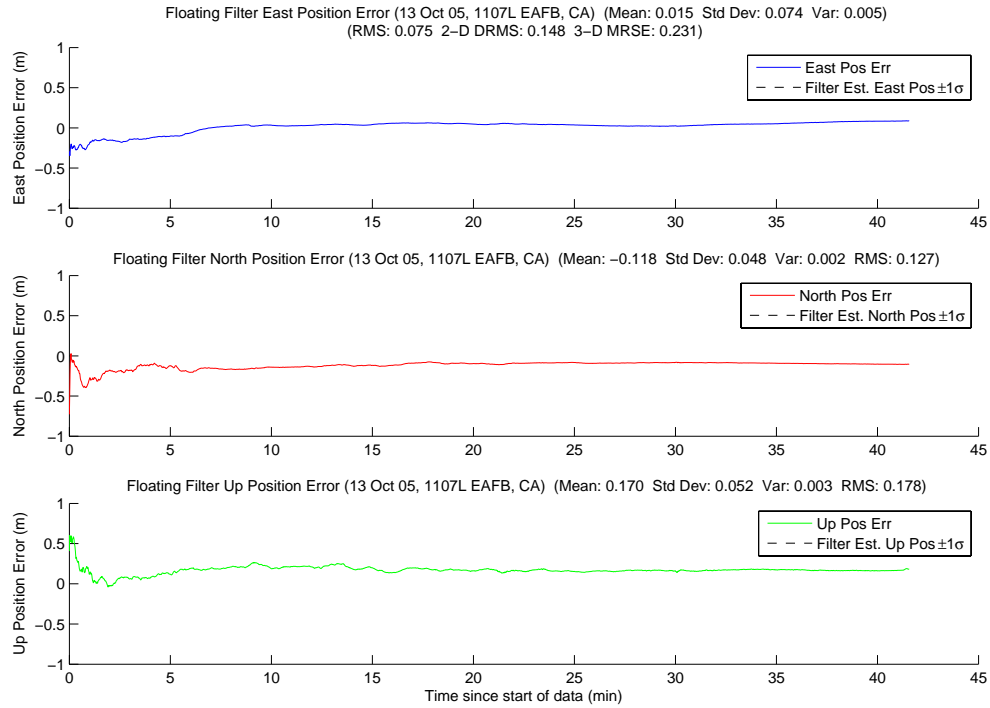


Figure A.342: Case F7.2: Floating Filter Position Errors

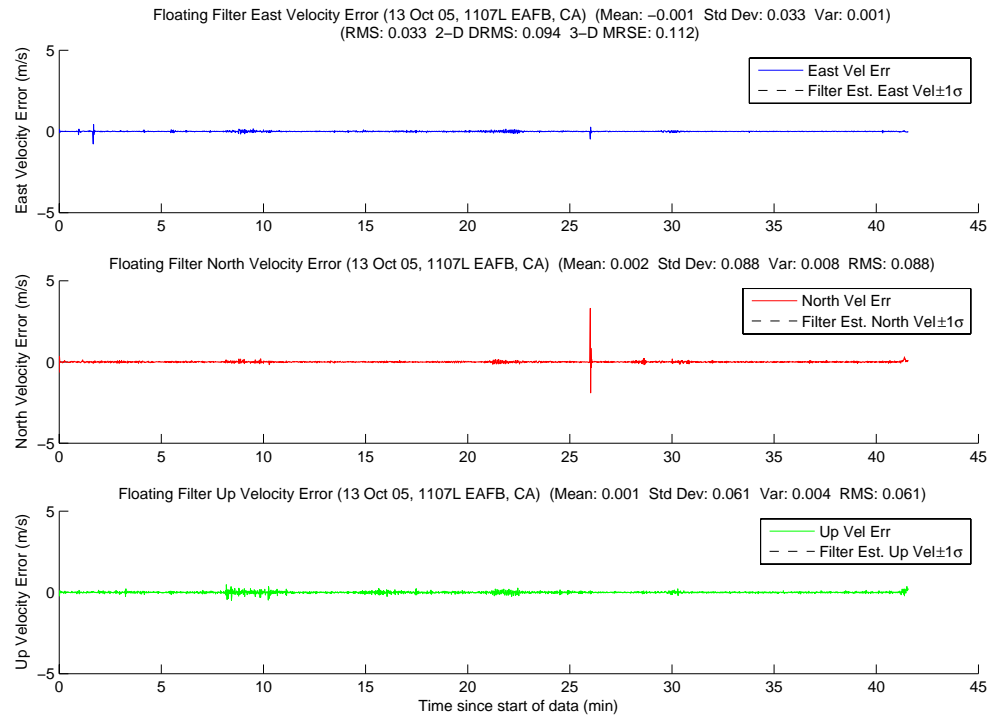


Figure A.343: Case F7.2: Floating Filter Velocity Errors

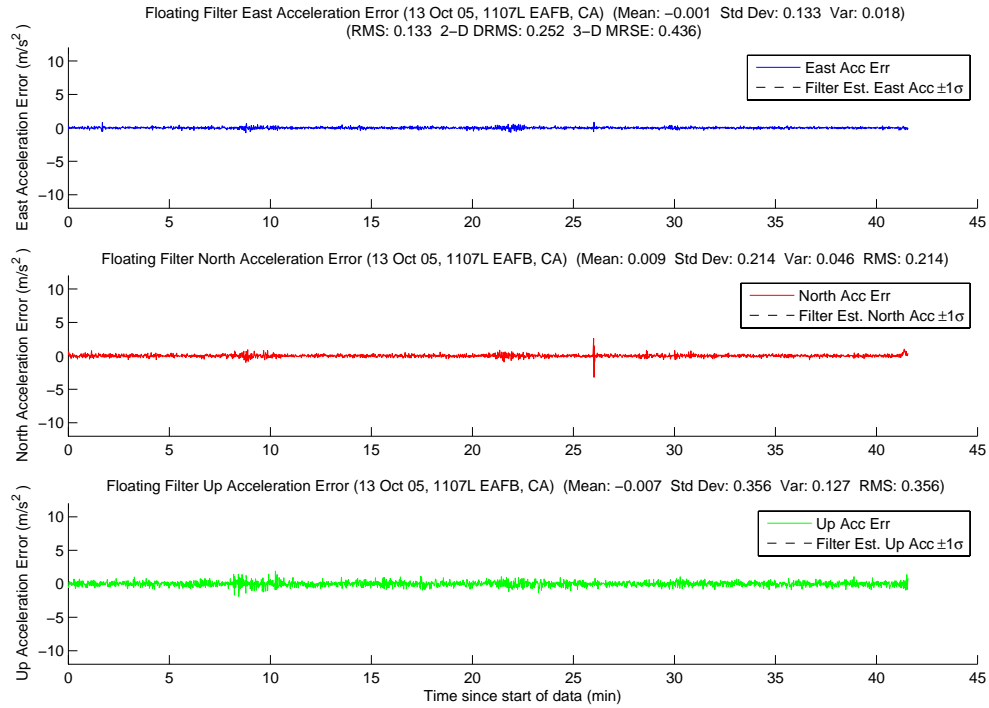


Figure A.344: Case F7.2: Floating Filter Acceleration Errors

Figure A.345 displays PRN 10 measurement information and the floating filter ambiguity estimate. The breaks in the floating filter's 1σ estimation at 9.5, 22, and 30 minutes were due to masking. The floating filter had trouble converging on the true ambiguity for PRN 10. This was due, in part, to the satellite's low elevation. However, the three resets were the largest contributors to the error (i.e. the filter did not have time converge, especially with a low elevated SV). The rest of the satellite's plots follow in Figures A.346, A.347, A.348, A.349, A.350, and A.351.

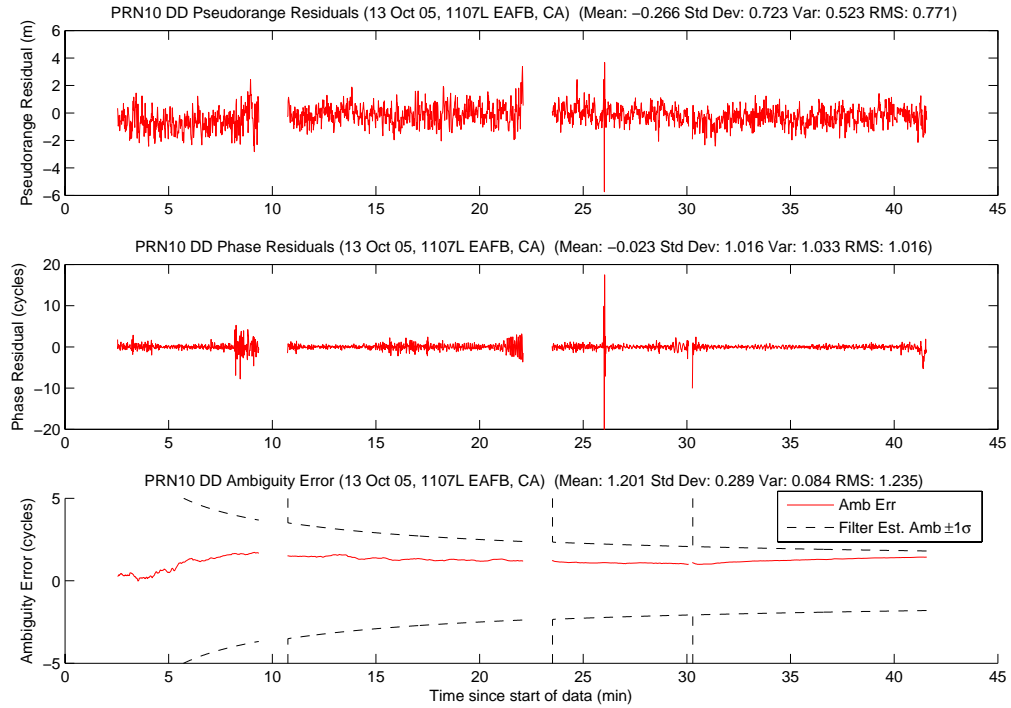


Figure A.345: Case F7.2: Satellite 10 Measurements

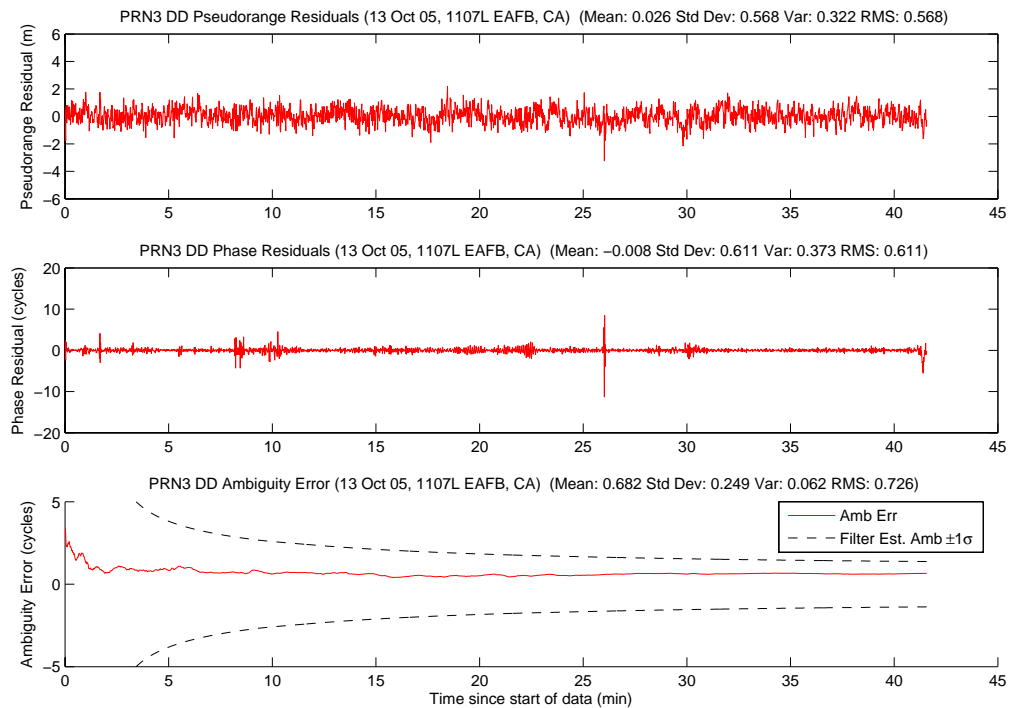


Figure A.346: Case F7.2: Satellite 3 Measurements

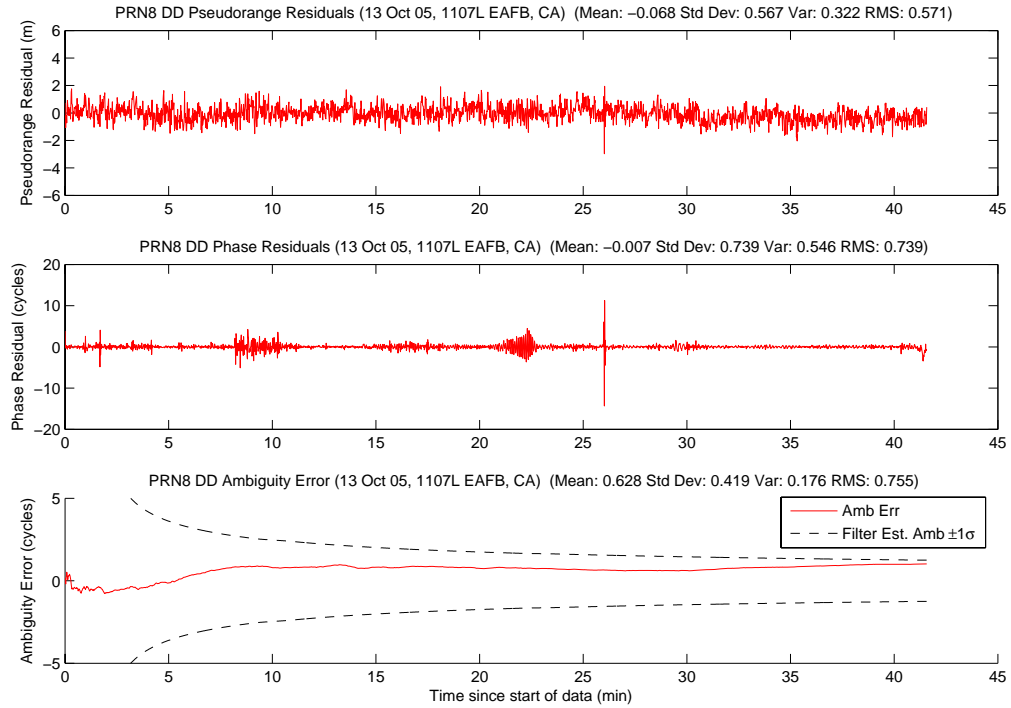


Figure A.347: Case F7.2: Satellite 8 Measurements

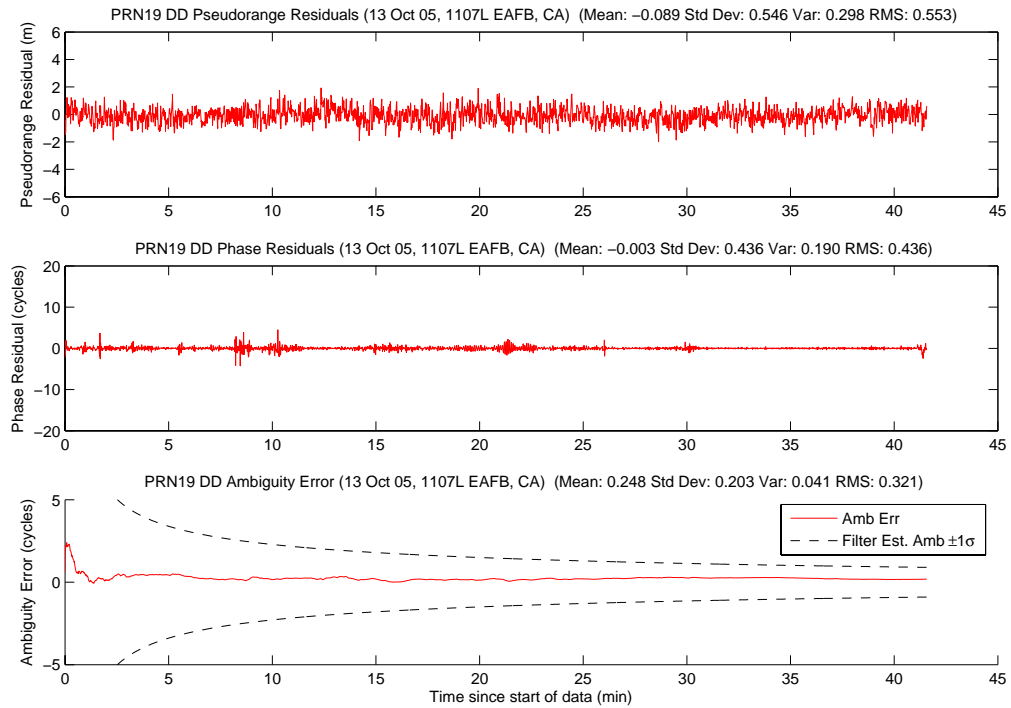


Figure A.348: Case F7.2: Satellite 19 Measurements

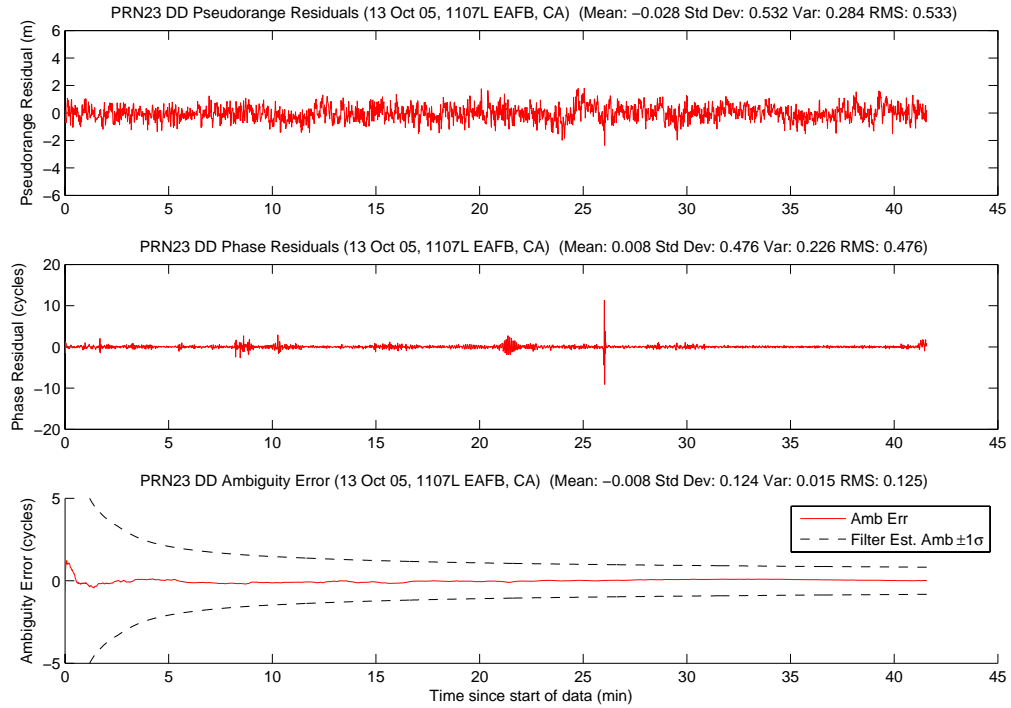


Figure A.349: Case F7.2: Satellite 23 Measurements

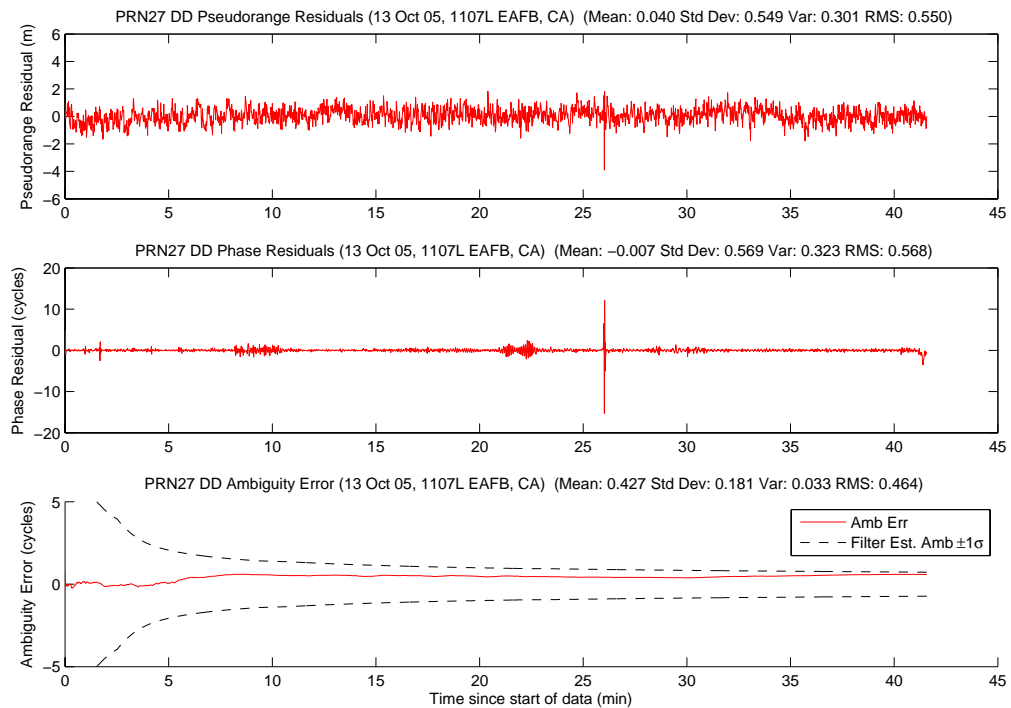


Figure A.350: Case F7.2: Satellite 27 Measurements

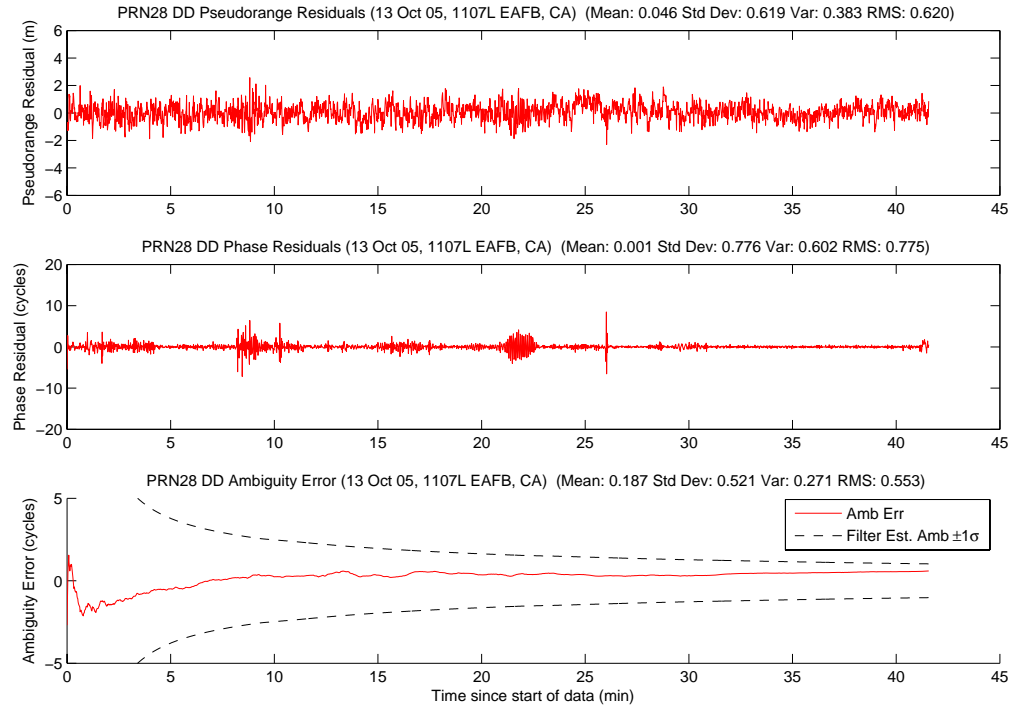


Figure A.351: Case F7.2: Satellite 28 Measurements

Figure A.352 shows the number of visible satellites. There were three satellite dropouts during the run, and all of the dropouts were due to masking of PRN 10. PRN 10 was the satellite that rose at two minutes. It was low on the horizon, and prone to being masked by the horizontal tail or lead aircraft.

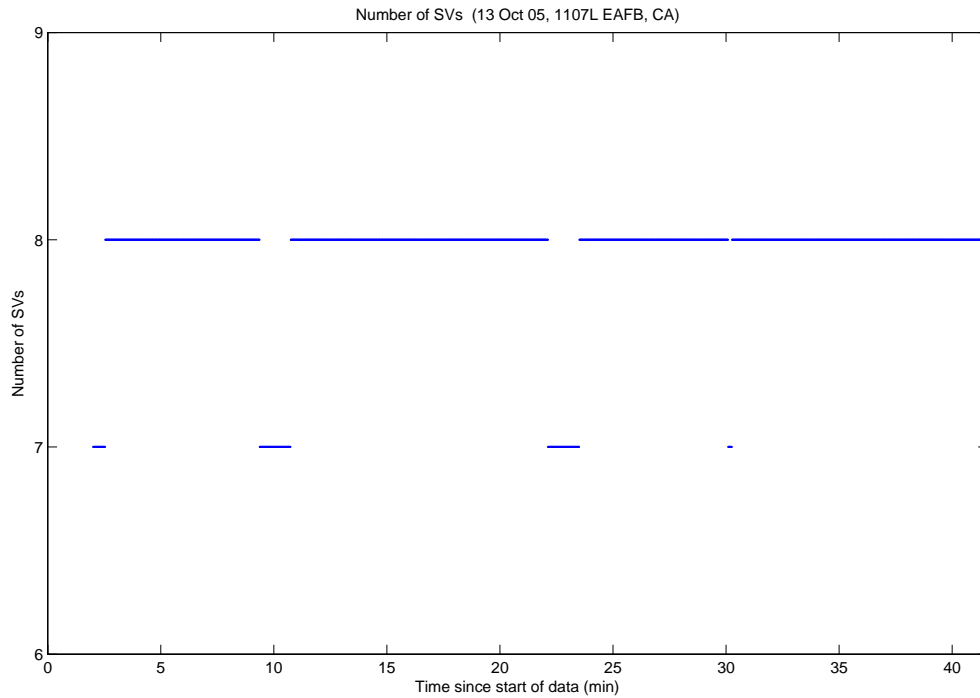


Figure A.352: Case F7.2: Visible Satellites

Figure A.353 shows all of the candidate ambiguity set's SOS residuals. Up until 22 minutes into the data run, there were numerous erroneous ambiguity set's residuals near the true set. After 22 minutes, the true ambiguity set is clearly visible which coincided with PRN 10 unmasking the second time. Figure A.354 displays the MMAE conditional probability. The probability plot is similar to several others with eight visible satellites. By three minutes, the true ambiguity set's probability was already 83 percent, and rising quickly.

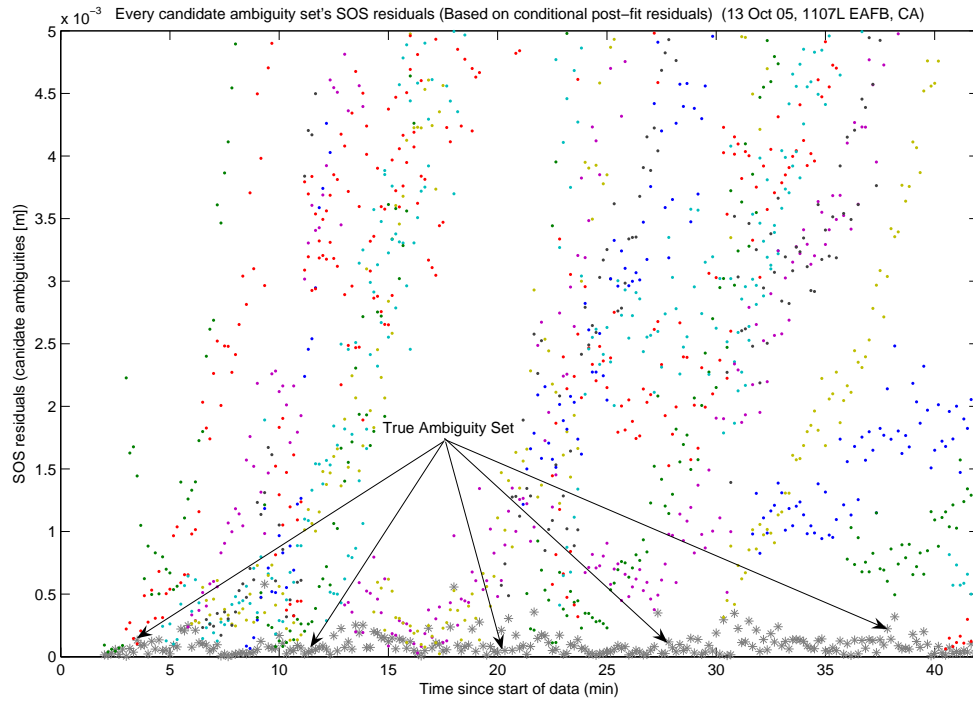


Figure A.353: Case F7.2: SOS Residuals

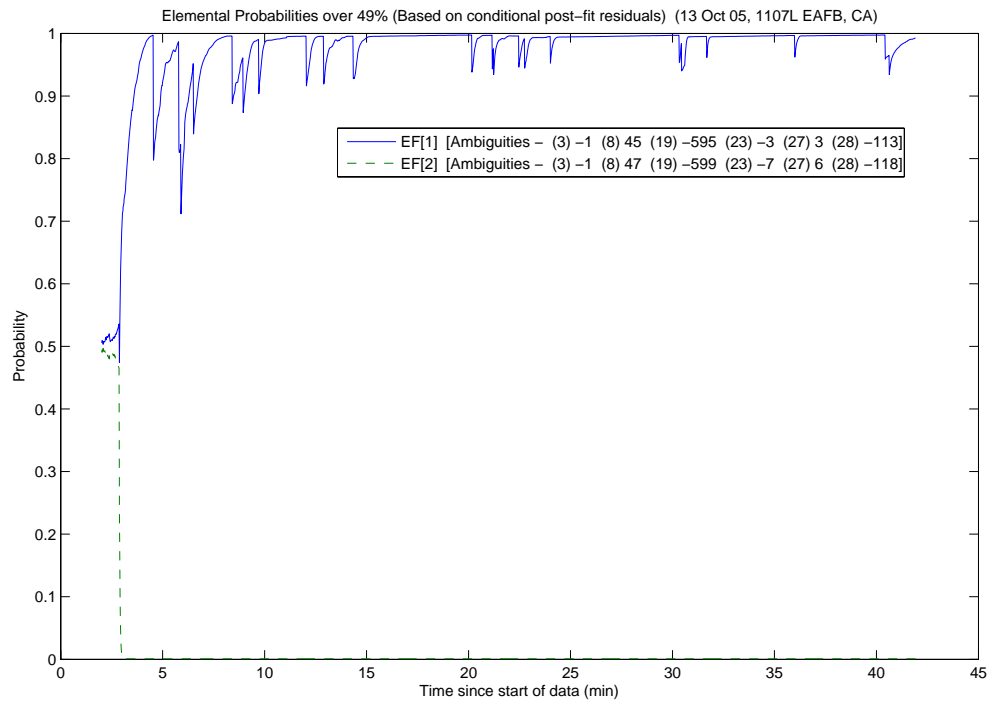


Figure A.354: Case F7.2: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.355 is EF[1], and Figure A.356 is EF[2].

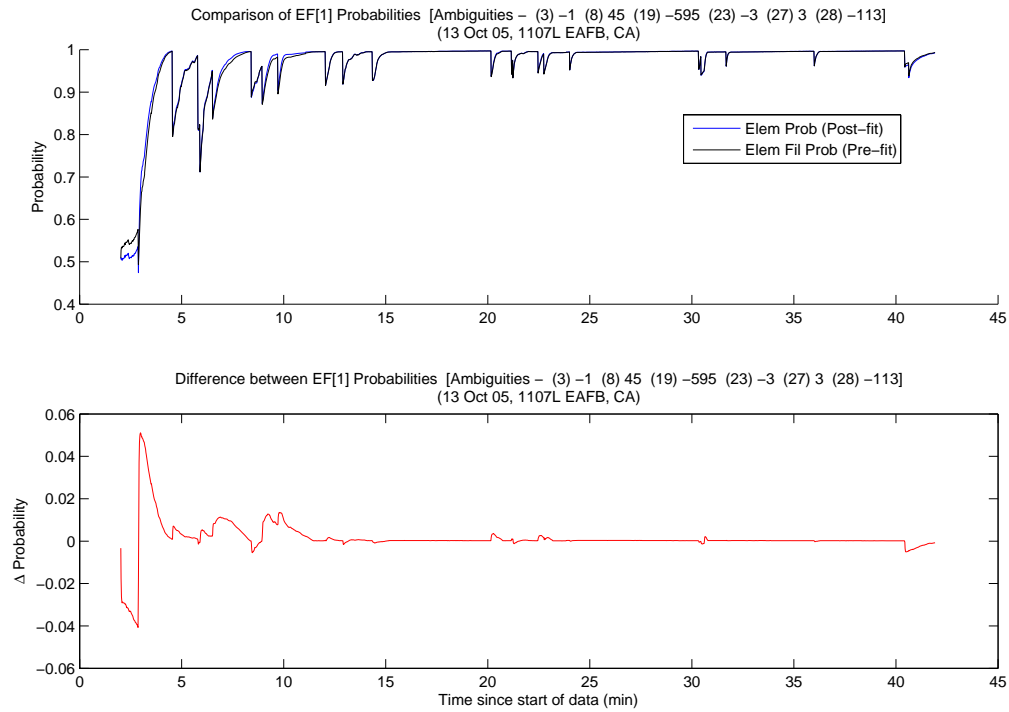


Figure A.355: Case F7.2: EF[1] Probability Comparison

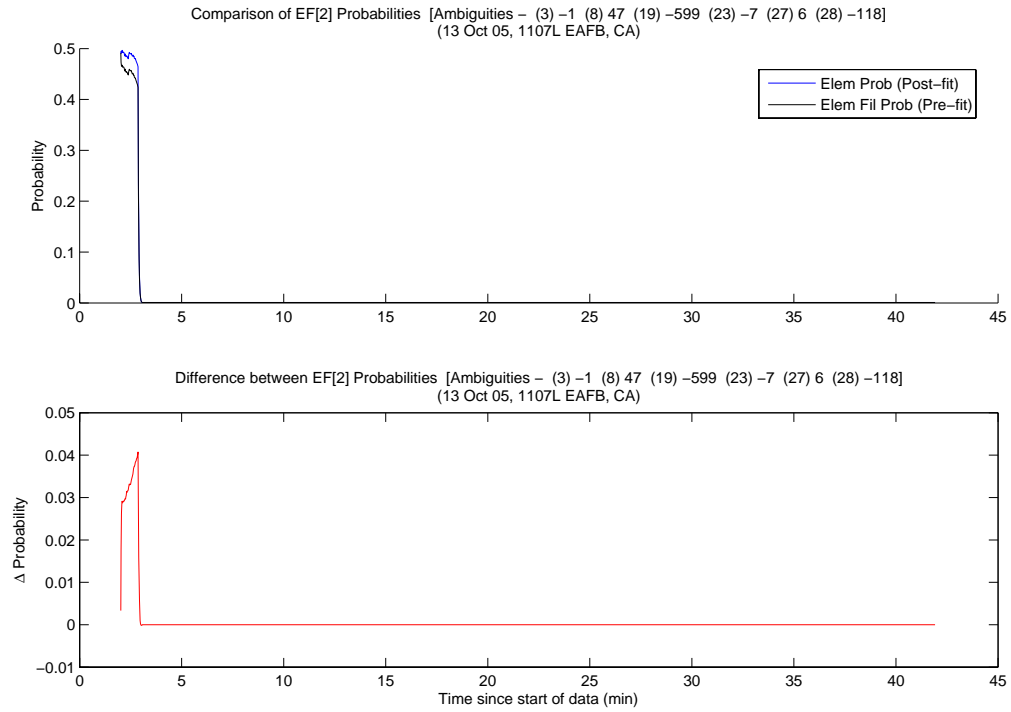


Figure A.356: Case F7.2: EF[2] Probability Comparison

The MMAE position error is shown in Figure A.357. The plot is similar to the last few data runs.

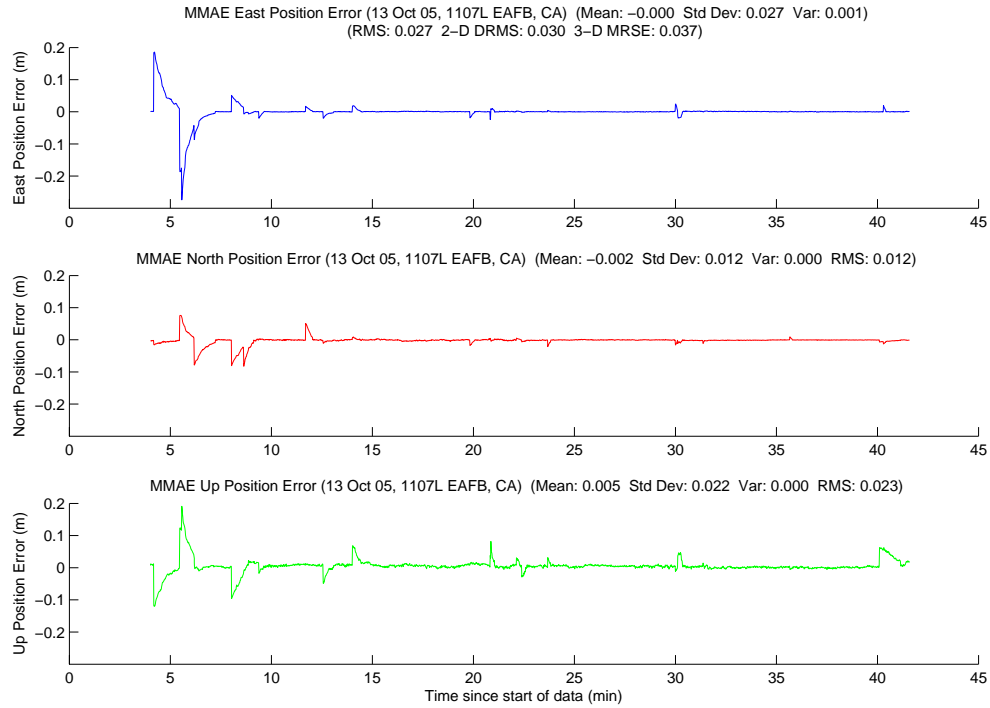


Figure A.357: Case F7.2: MMAE Position Error

Finally, Figure A.358 displays the position error plot for the minimum indicator variable. For the fourth straight data run, the minimum indicator MRSE was within millimeters. Lastly, Table A.19 gives a summary of each method's position error.

Table A.19: Case F7.2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	0.015	0.074	-0.118	0.048	0.170	0.052	0.148	0.231
MMAE	-0.000	0.027	-0.002	0.012	0.005	0.022	0.030	0.037
Min. Ind.	0.000	0.001	-0.001	0.001	0.005	0.004	0.002	0.007

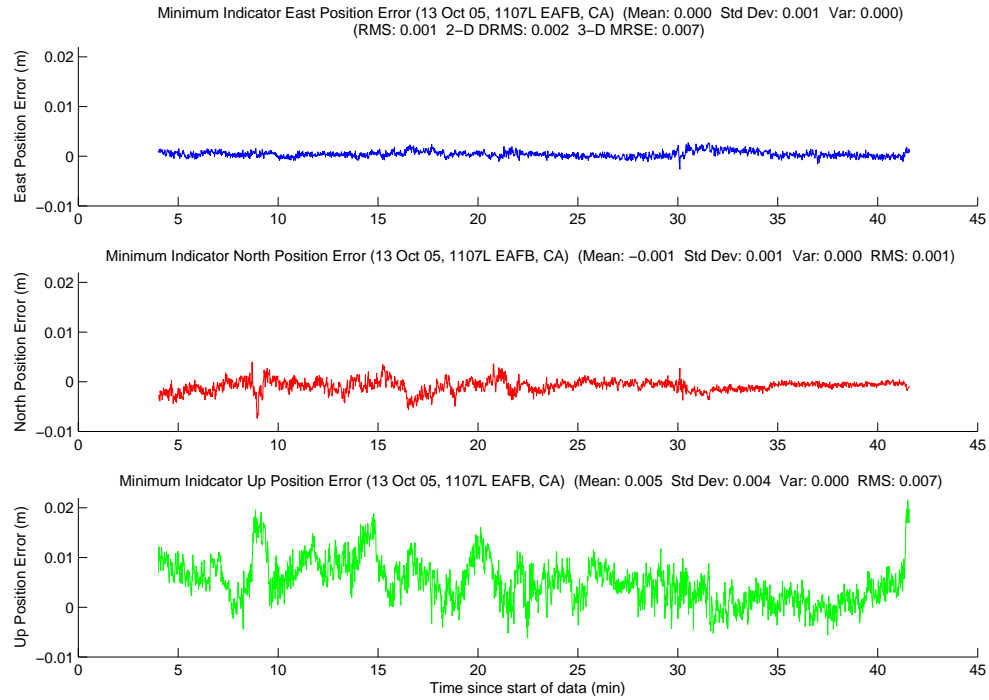


Figure A.358: Case F7.2: Minimum Indicator Position Error

A.2.16 Case F8.1: Flight Test 8, First Data Set . The last flight test was flown on 14 October 2005. The first data run started at 0901L and lasted 59 minutes, 33 seconds. The weather was clear with light winds, and the aircraft ranged from 11.0 to 125.8 meters apart during the run. Figure A.359 shows the minimum indicator North, East, Down relative positions.

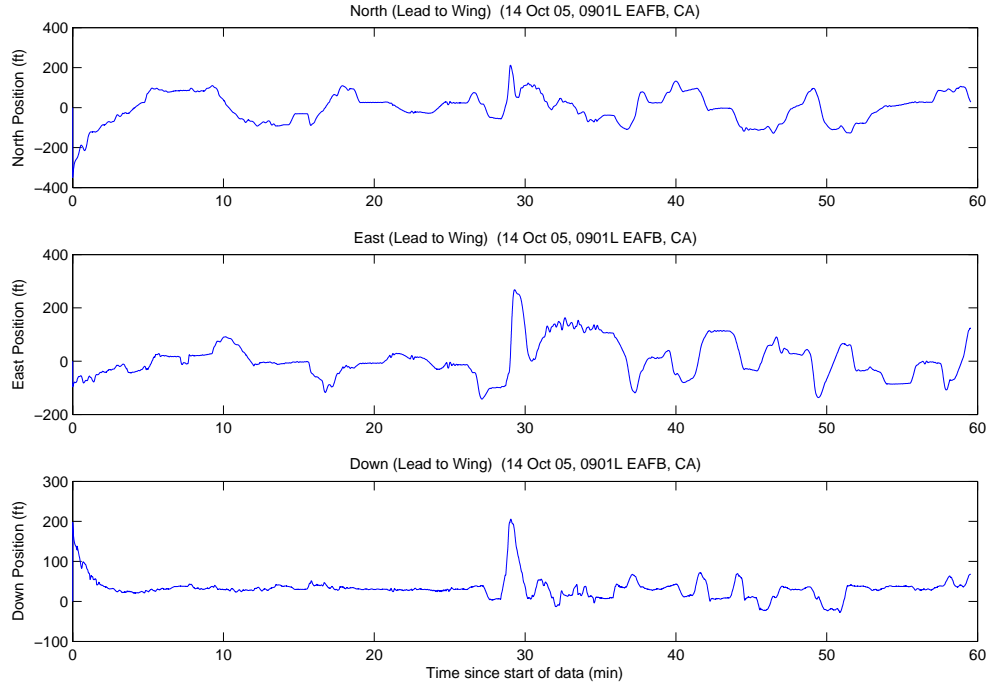


Figure A.359: Case F8.1: North, East, Down Relative Positions

A total of 29 unique ambiguity sets were returned from the LAMBDA function during the data run. The aircraft were 125.8 meters apart when the system started, with nine visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 3390 of 3451 epochs (98.2 percent), with a maximum ratio of 303.9.

Figures A.360, A.361, and A.362 show the floating filter position errors. The floating filter performed as expected, and was very similar to the previous data runs. The floating filter's 1σ estimate of the error is outside the plot scales in all three figures.

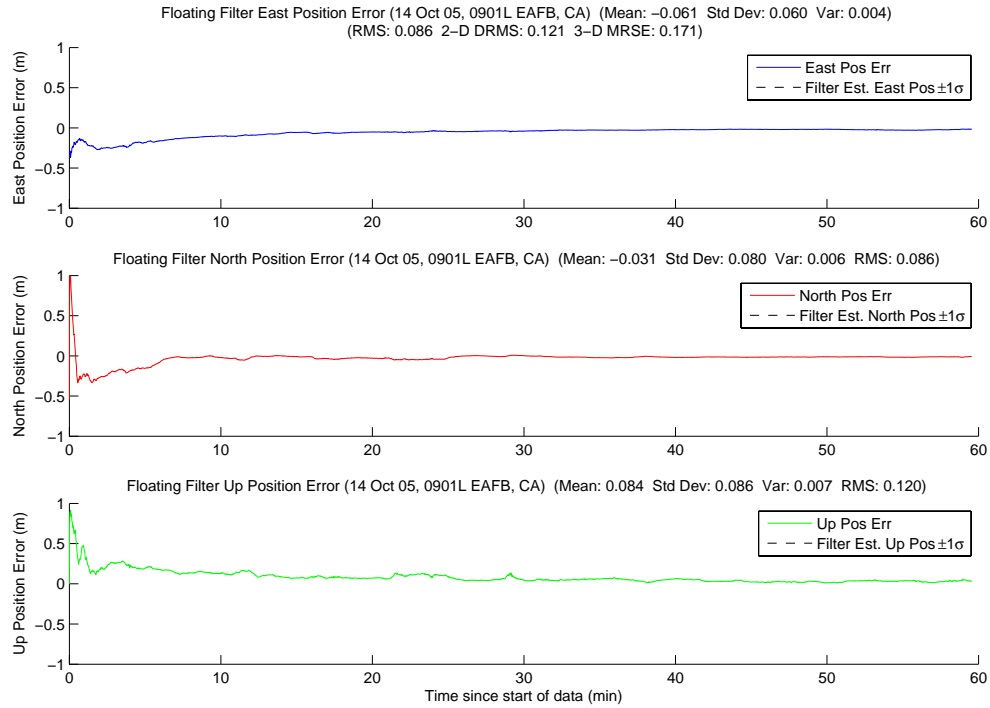


Figure A.360: Case F8.1: Floating Filter Position Errors

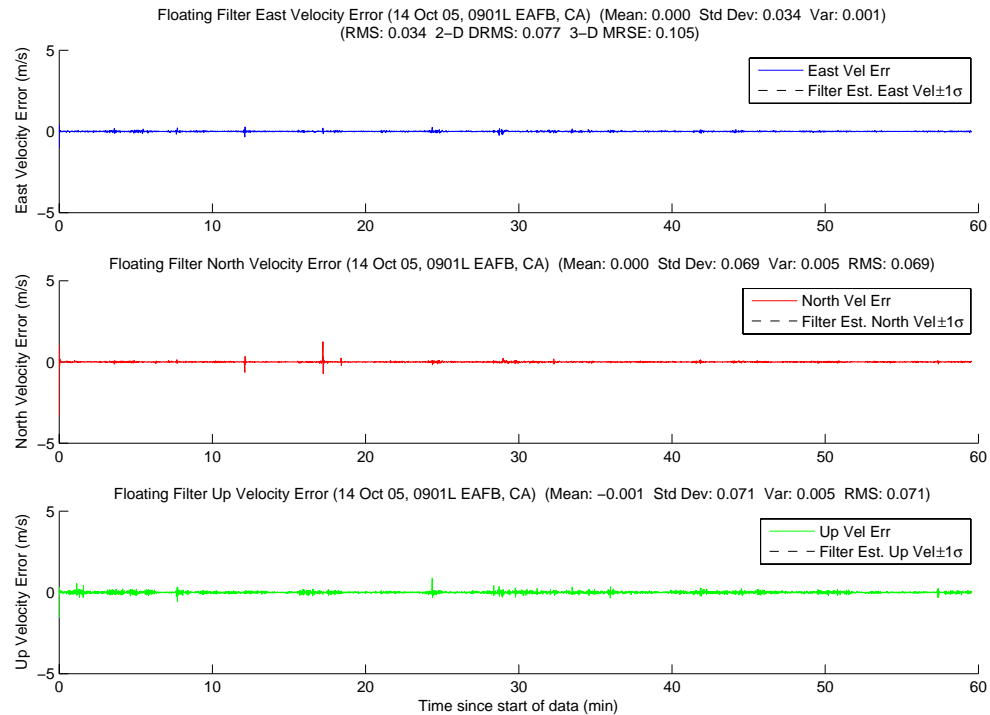


Figure A.361: Case F8.1: Floating Filter Velocity Errors

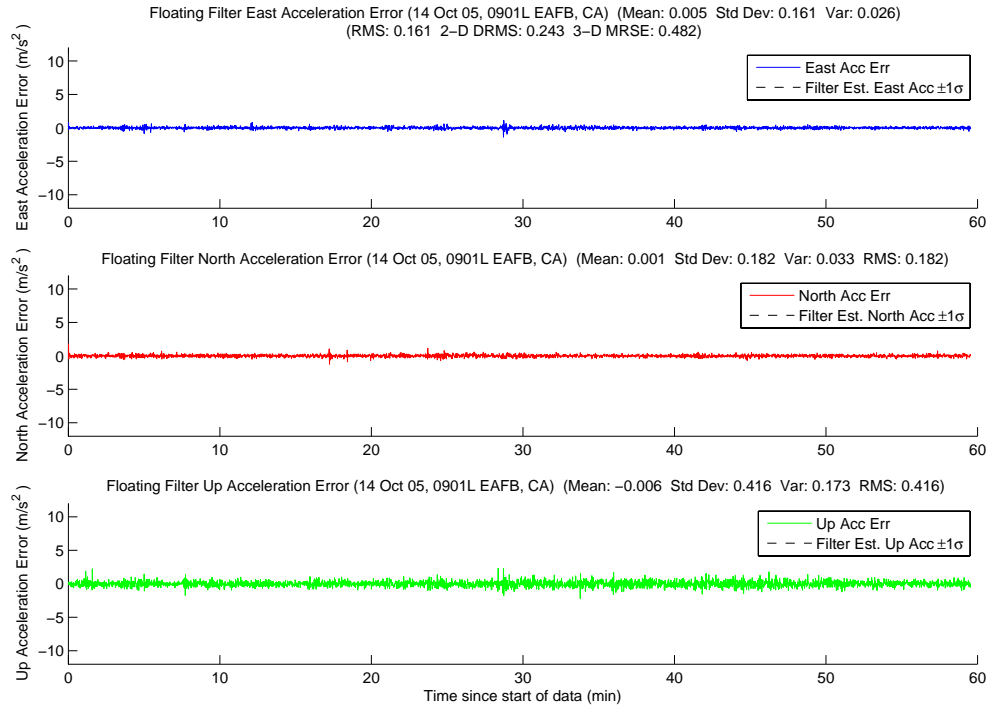


Figure A.362: Case F8.1: Floating Filter Acceleration Errors

Figure A.363 displays PRN 4 measurement information and the floating filter ambiguity estimate. The breaks in the floating filter's 1σ estimation at 10.5 and 24 minutes were due to masking, and the satellite set at 25.5 minutes. The rest of the satellite's plots follow in Figures A.364, A.365, A.366, A.367, A.368, A.369, A.370, A.371, A.372, and A.373.

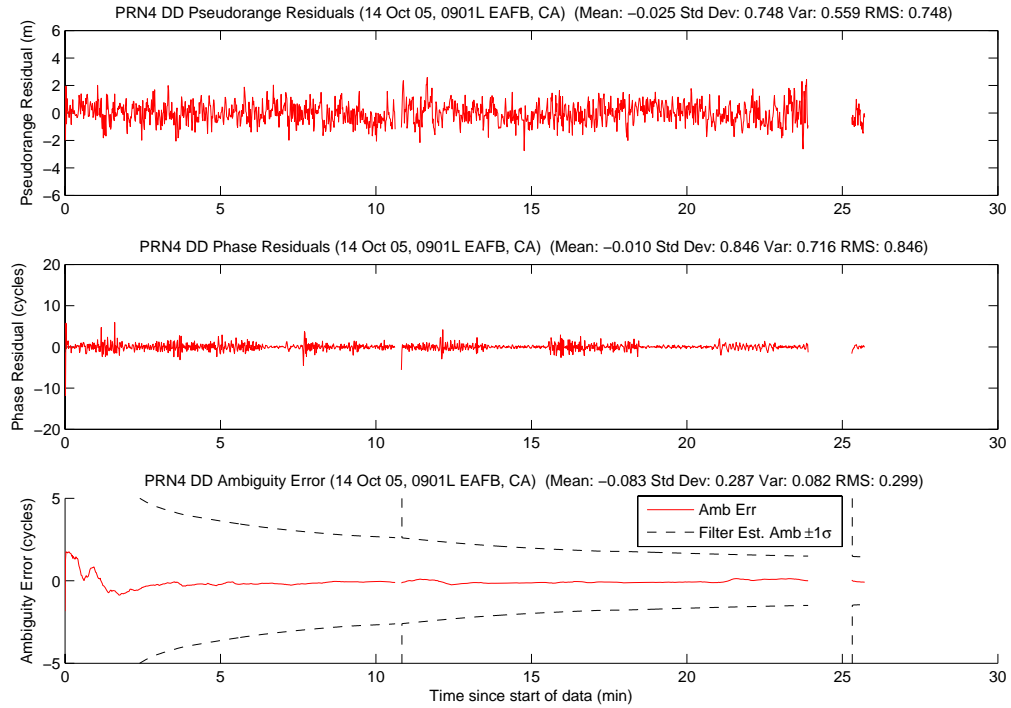


Figure A.363: Case F8.1: Satellite 4 Measurements

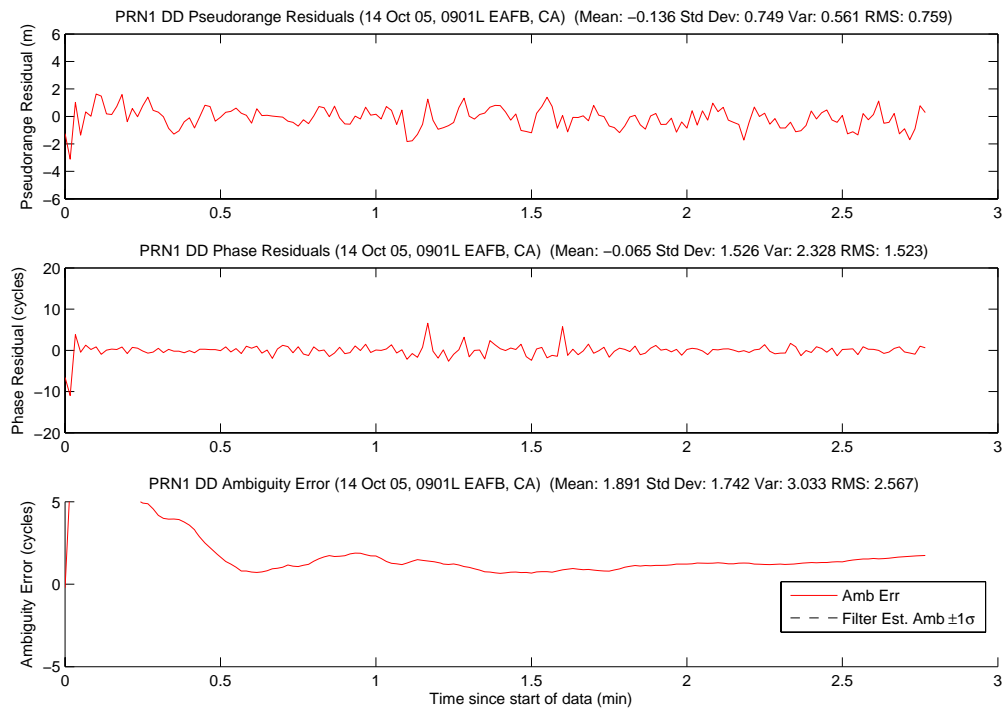


Figure A.364: Case F8.1: Satellite 1 Measurements

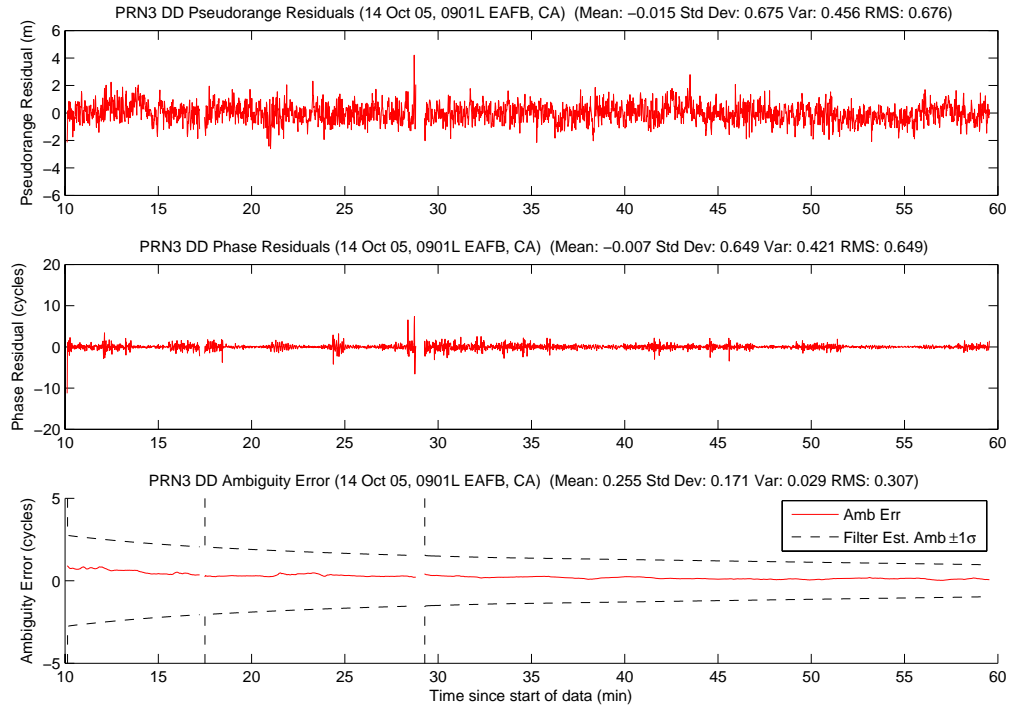


Figure A.365: Case F8.1: Satellite 3 Measurements

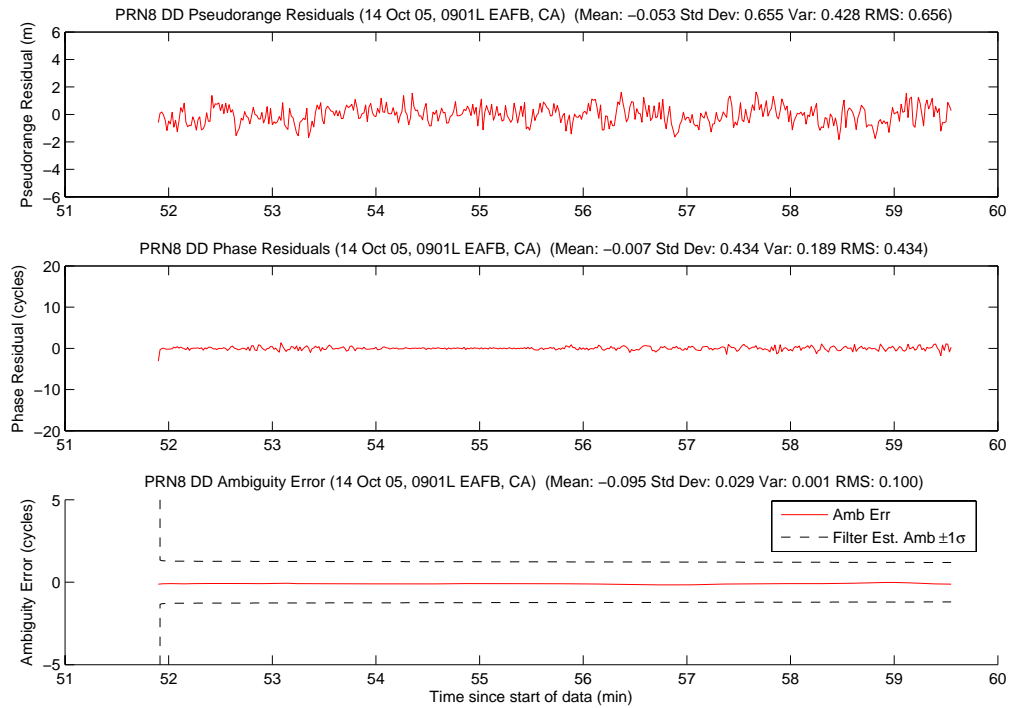


Figure A.366: Case F8.1: Satellite 8 Measurements

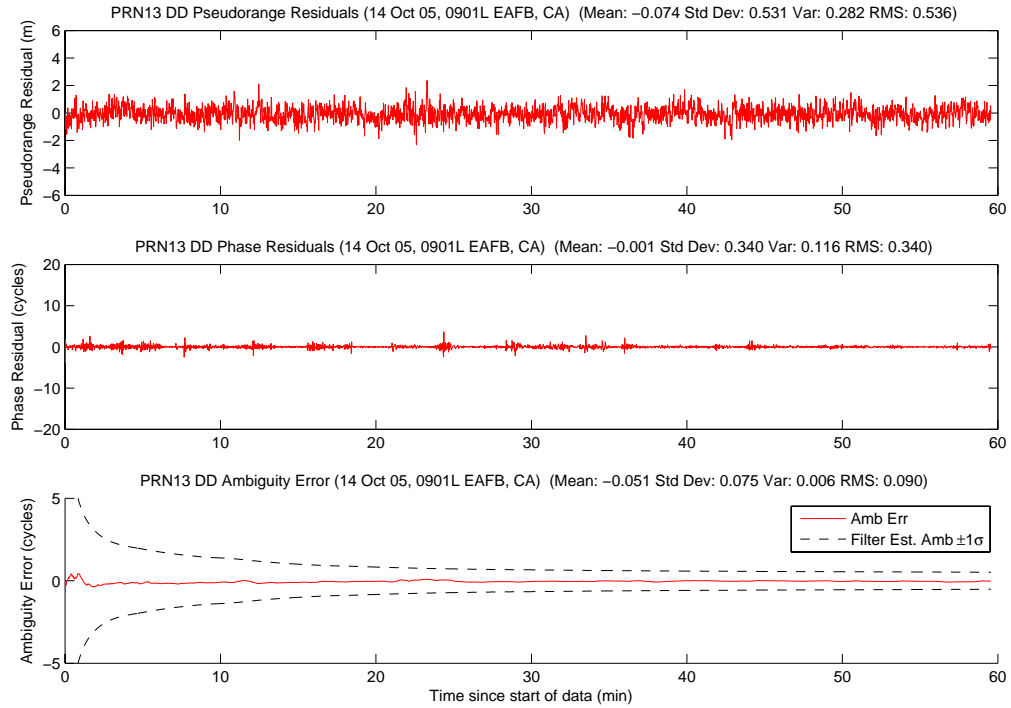


Figure A.367: Case F8.1: Satellite 13 Measurements

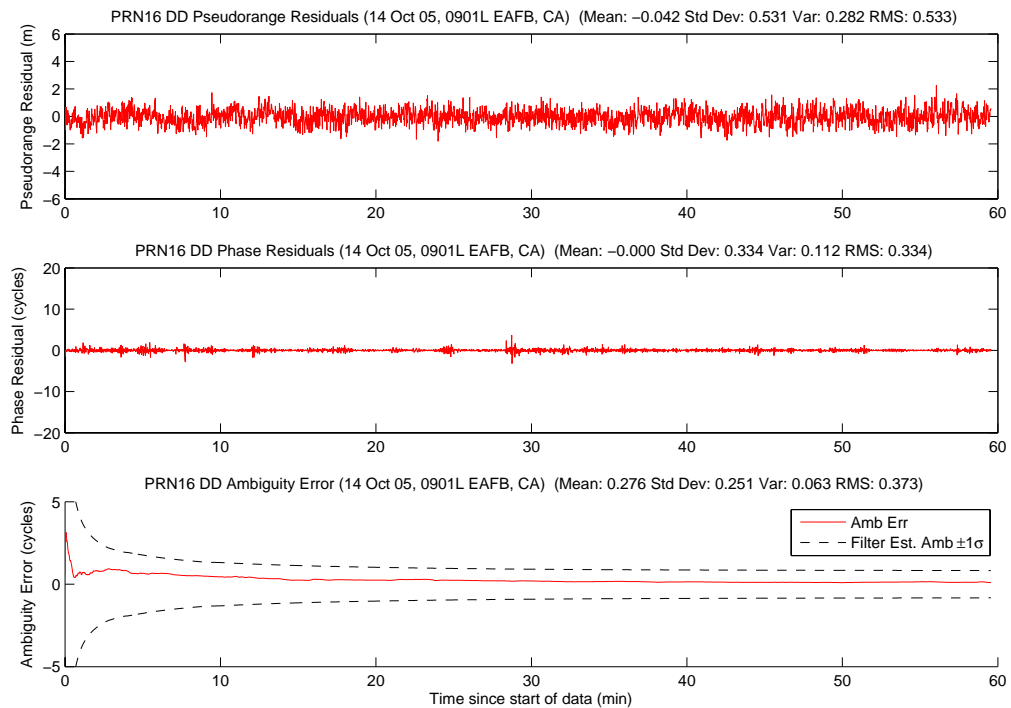


Figure A.368: Case F8.1: Satellite 16 Measurements

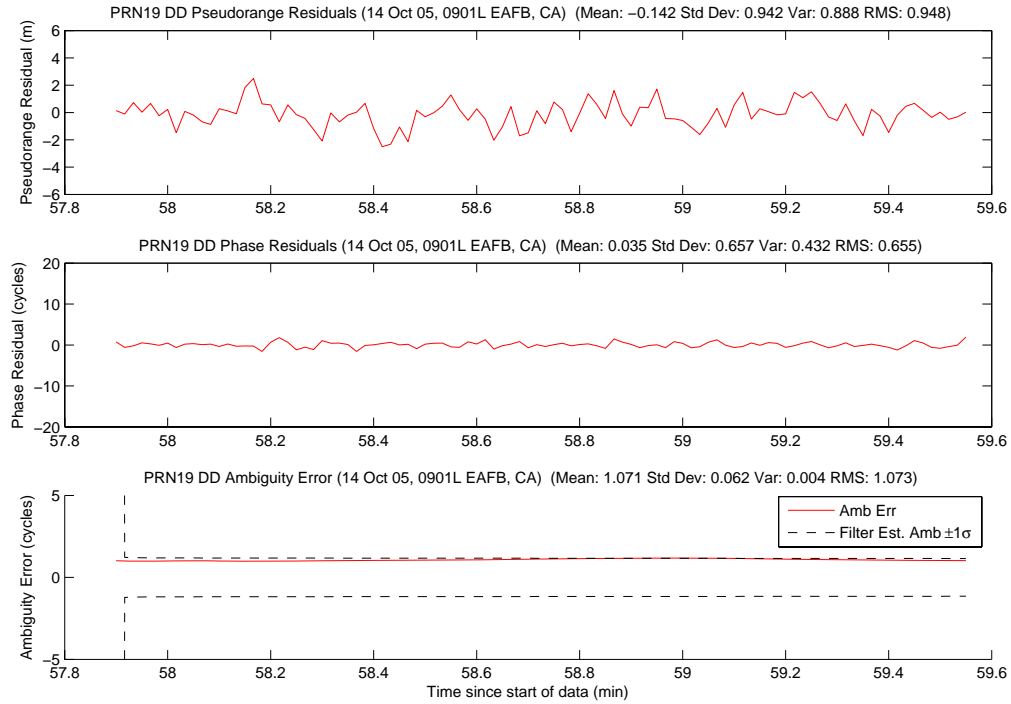


Figure A.369: Case F8.1: Satellite 19 Measurements

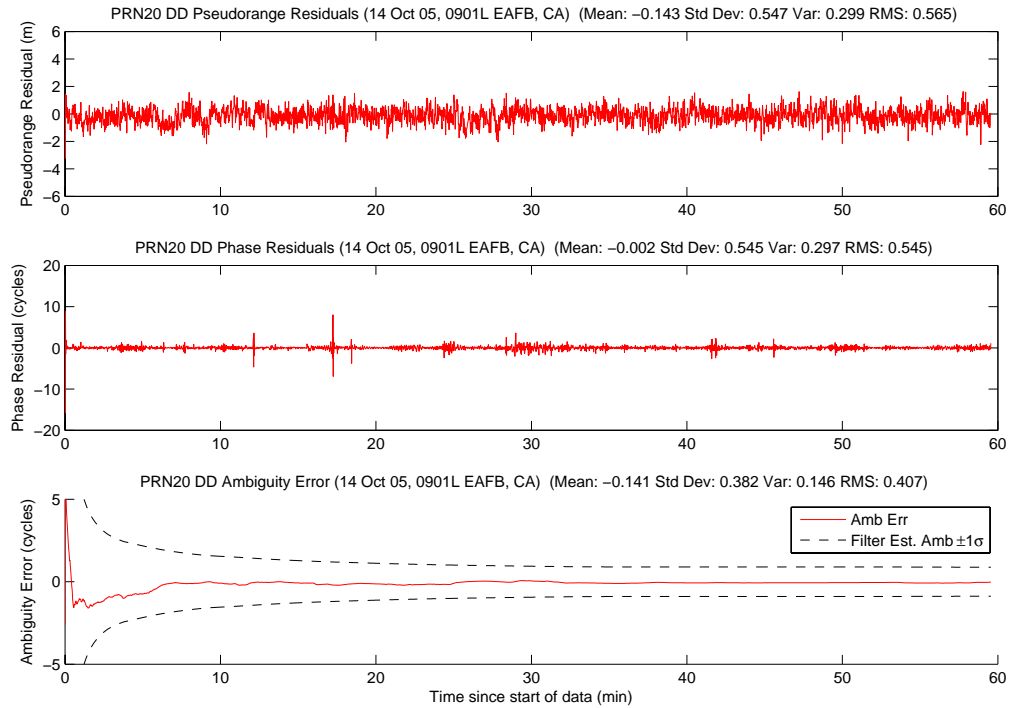


Figure A.370: Case F8.1: Satellite 20 Measurements

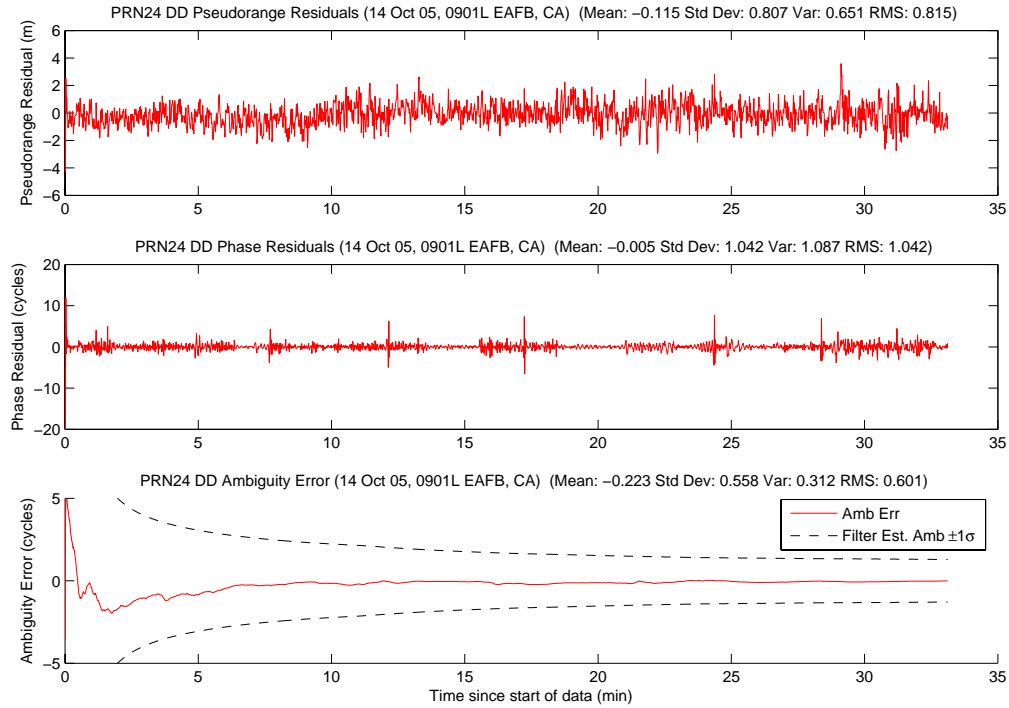


Figure A.371: Case F8.1: Satellite 24 Measurements

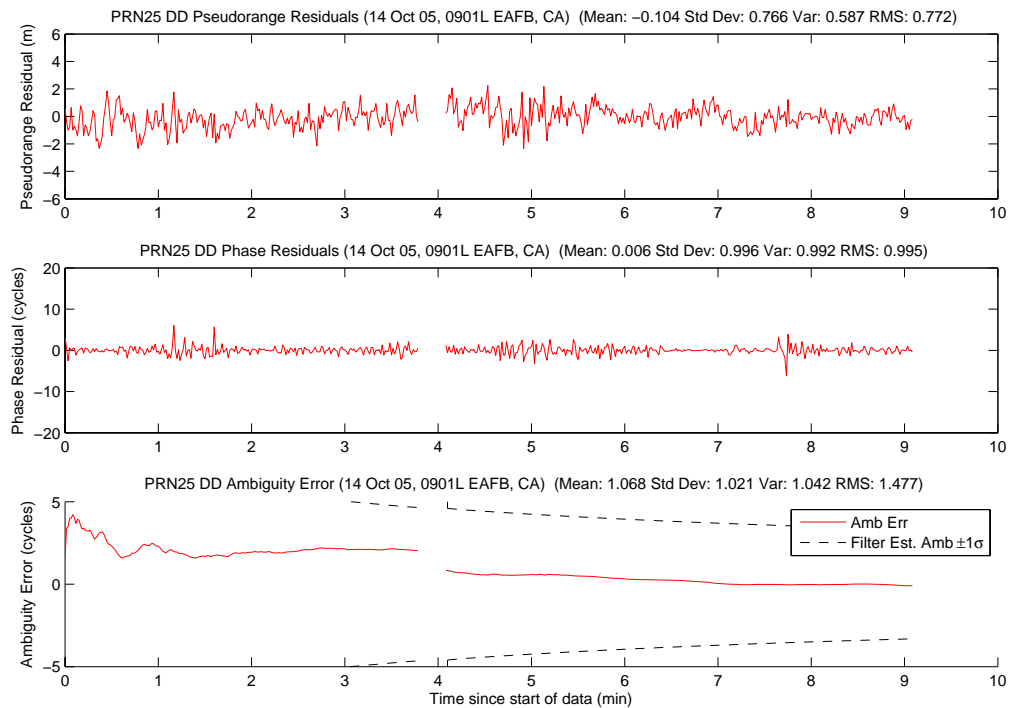


Figure A.372: Case F8.1: Satellite 25 Measurements

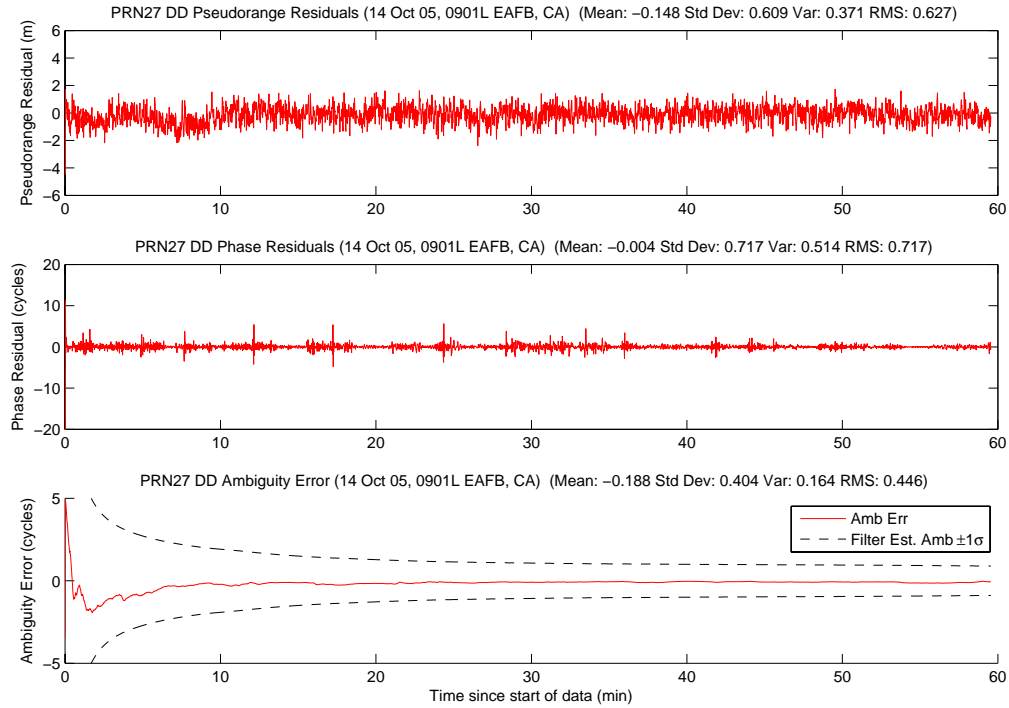


Figure A.373: Case F8.1: Satellite 27 Measurements

Figure A.374 shows the number of visible satellites. There was a lot of satellite activity during this data run. Five dropouts occurred, and seven satellites rose/set, the most of any data run. PRNs 1 and 24 set at 3 and 33 minutes, respectively. PRN 25 masked at 5 minutes, and then set at 9 minutes. PRN 4 masked at 10.5 and 24 minutes, and then set at 25.5 minutes. PRNs 8 and 19 rose at 52 and 58 minutes, respectively. Finally, PRN 3 rose at 10 minutes and then masked at 17 and 29 minutes. Even with all of the low elevated satellites and maskings, the filter, MMAE, and minimum indicator performed well.

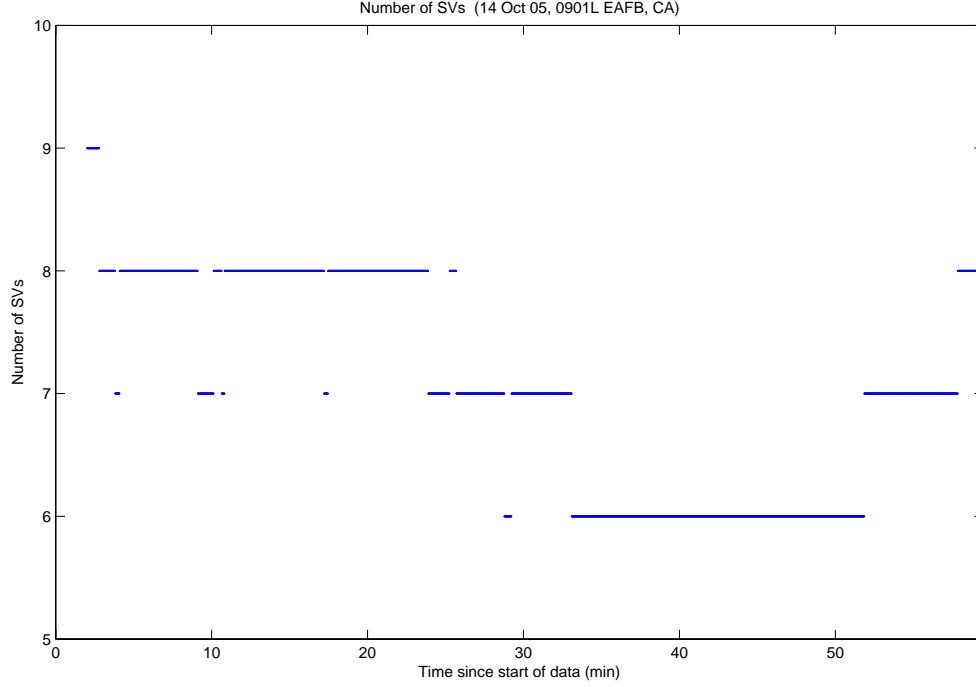


Figure A.374: Case F8.1: Visible Satellites

Figure A.375 shows all of the candidate ambiguity set's SOS residuals. At approximately 32 minutes, when there were only six visible satellites, the true ambiguity set's residuals would be unidentifiable without emphasis. At 52 minutes, when PRN 8 rose, the true ambiguity set's residuals became easy to recognize. To efficiently determine the true ambiguity set with the algorithm developed in this thesis, a minimum of six visible satellites were required. Figure A.376 displays the MMAE conditional probability. The large dip in the true ambiguity set's probability, starting at 11 minutes, was due to an erroneous ambiguity set's residuals going below the true set. The erroneous set is identified by open blue squares in Figure A.375. The other small dip in the true set's probability from 25 to 30 minutes was due to another erroneous set, represented by red dots on the residual plot. After 30 minutes, there were only six visible satellites. That was the reason there were numerous erroneous ambiguity set's residuals close to the true set, which caused the small probability dips from 30 minutes until 40 minutes.

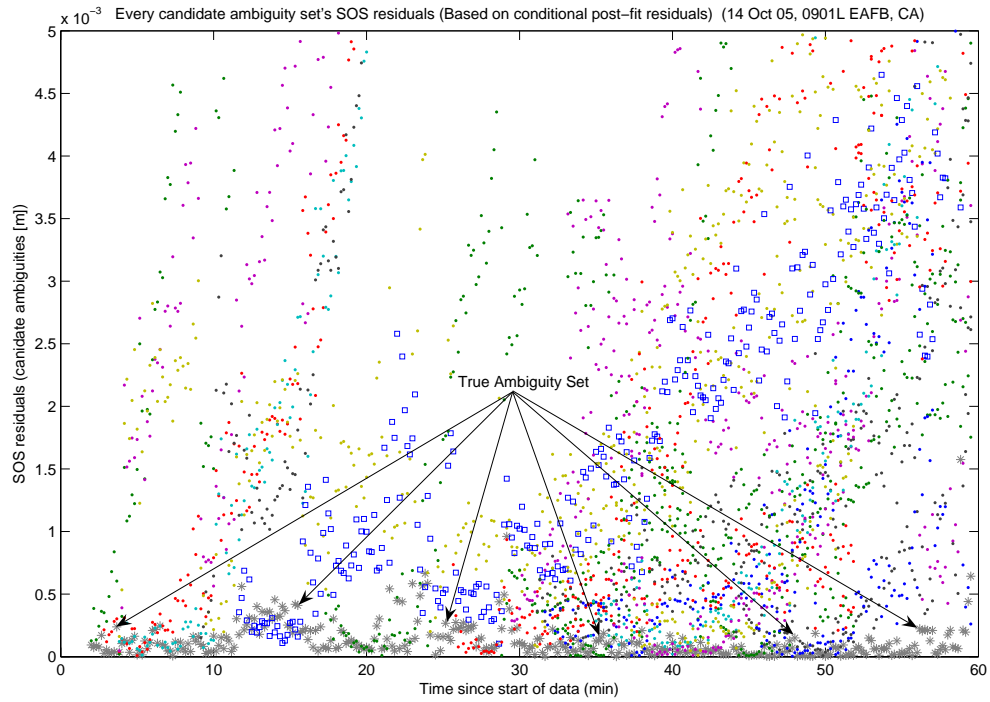


Figure A.375: Case F8.1: SOS Residuals

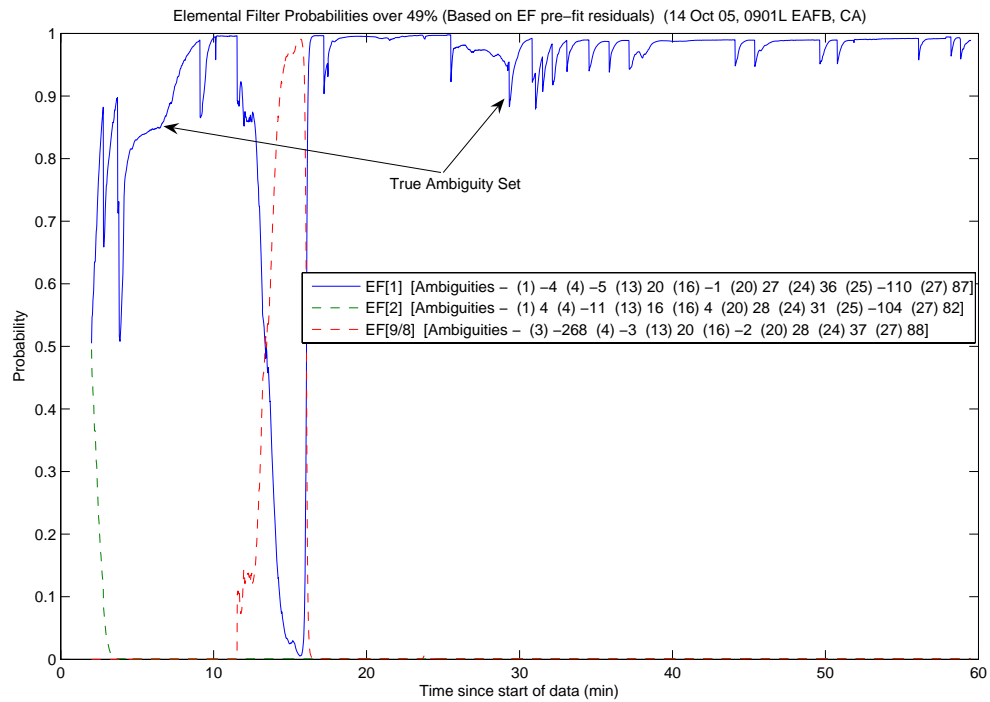


Figure A.376: Case F8.1: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.377 is EF[1], Figure A.378 is EF[2], and Figure A.379 is EF[9].

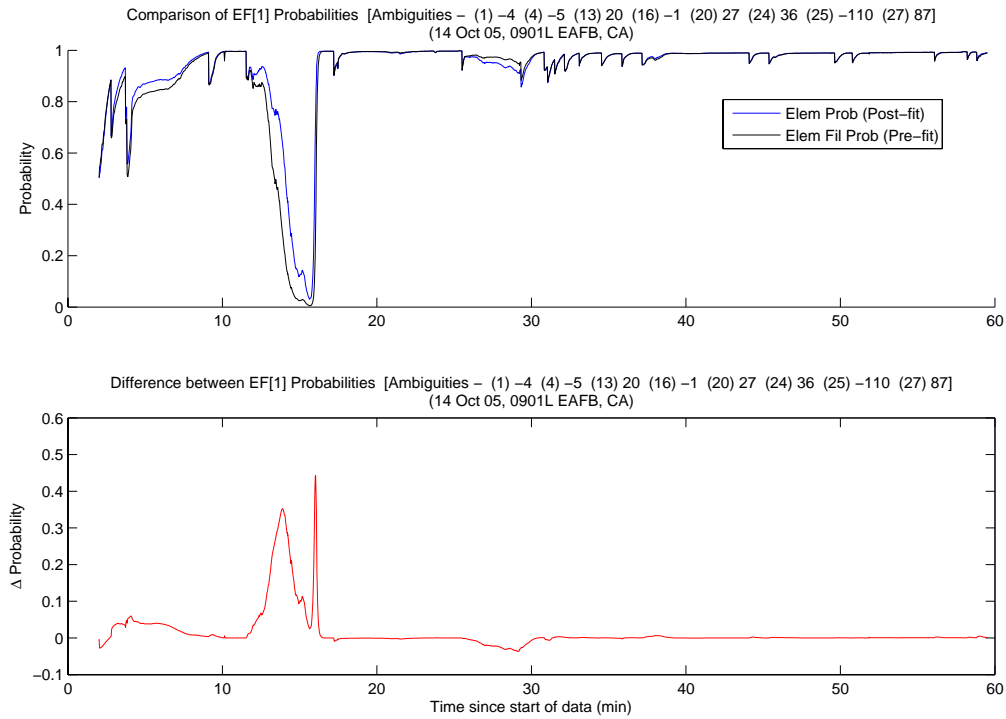


Figure A.377: Case F8.1: EF[1] Probability Comparison

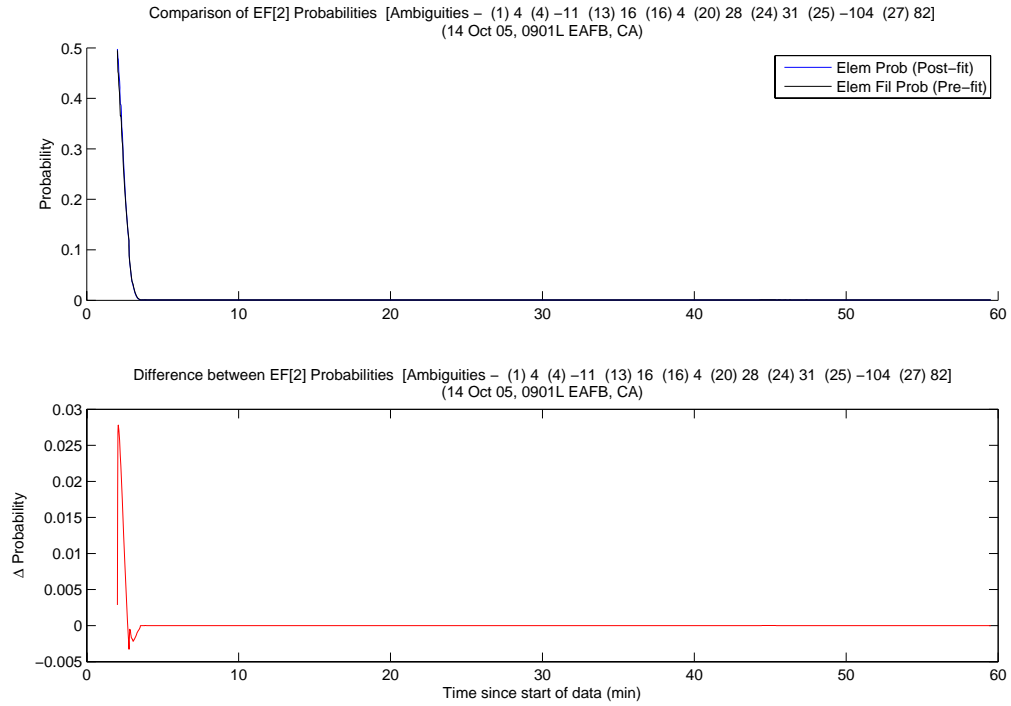


Figure A.378: Case F8.1: EF[2] Probability Comparison

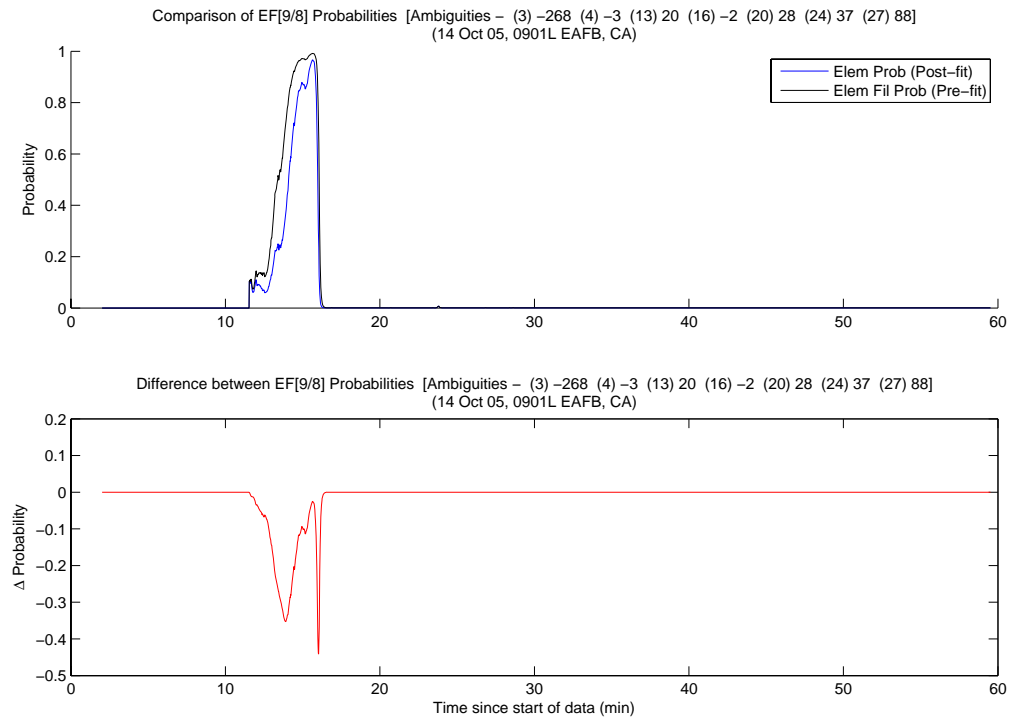


Figure A.379: Case F8.1: EF[9] Probability Comparison

The MMAE position error is shown in Figure A.380. The one large spike correlated to the large dip in the probability plot.

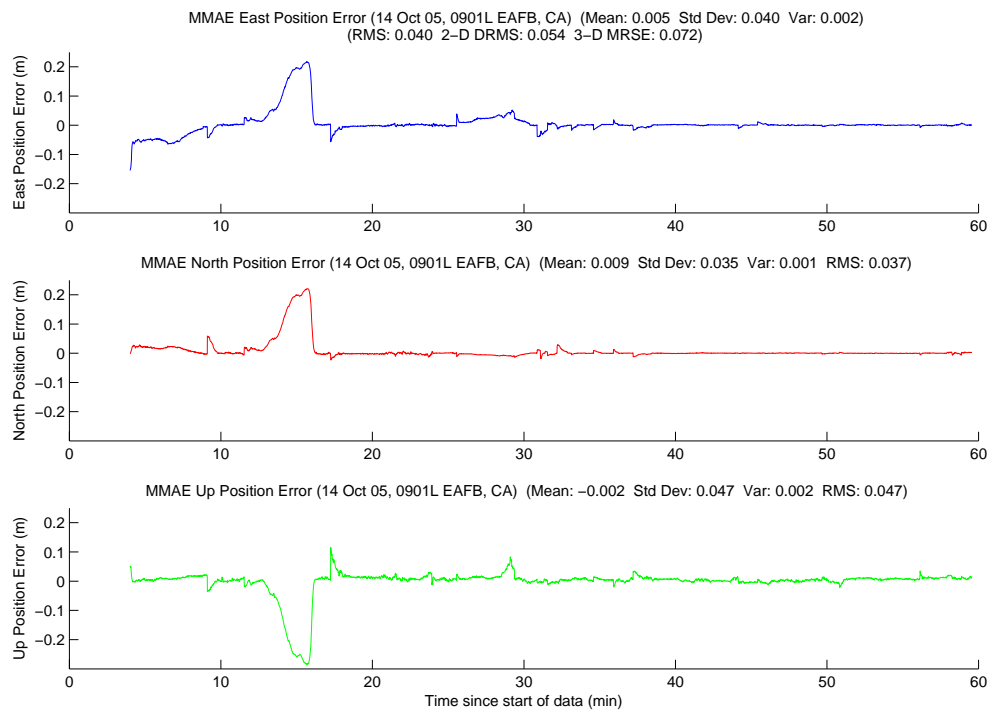


Figure A.380: Case F8.1: MMAE Position Error

Finally, Figure A.381 displays the position error plot for the minimum indicator variable. The minimum indicator performed great. The spike in the Up error at 29 minutes corresponded to PRN 3 masking, a low elevated satellite. Lastly, Table A.20 gives a summary of each method's position error.

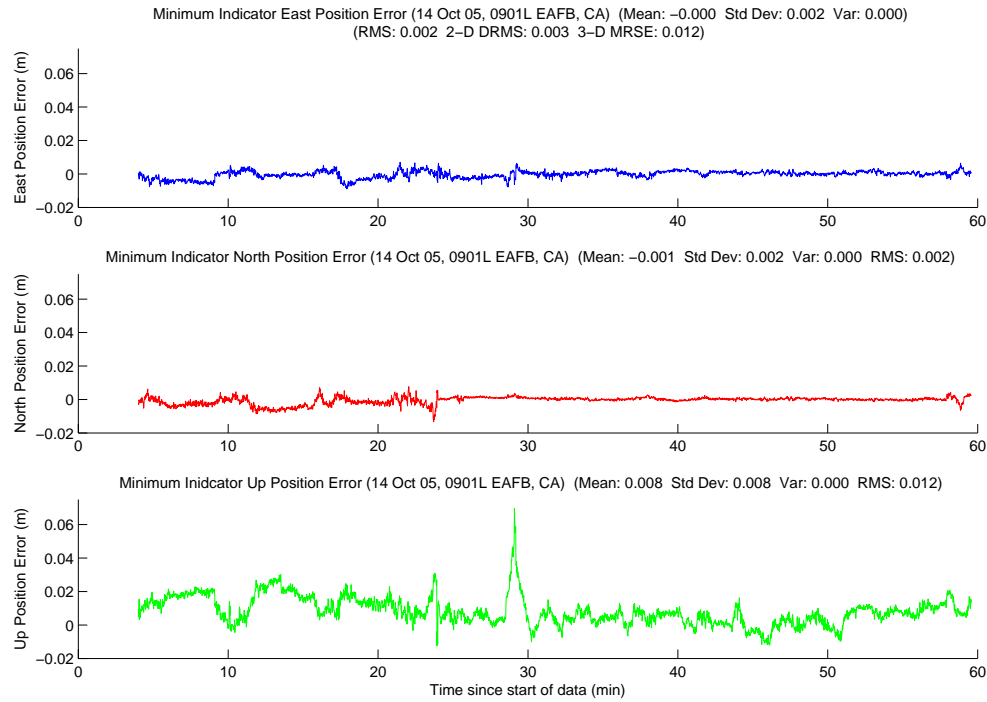


Figure A.381: Case F8.1: Minimum Indicator Position Error

Table A.20: Case F8.1: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.061	0.060	-0.031	0.080	0.084	0.086	0.121	0.171
MMAE	0.005	0.040	0.009	0.035	-0.002	0.047	0.054	0.072
Min. Ind.	-0.000	0.002	-0.001	0.002	0.008	0.008	0.003	0.012

A.2.17 *Case F8.2: Flight Test 8, Second Data Set* . Case F8.2 was the final flight test data run. The run started at 1001L and lasted 39 minutes, 28 seconds. The aircraft ranged from 10.9 to 633.5 meters apart during the run. Figure A.382 shows the minimum indicator North, East, Down relative positions.

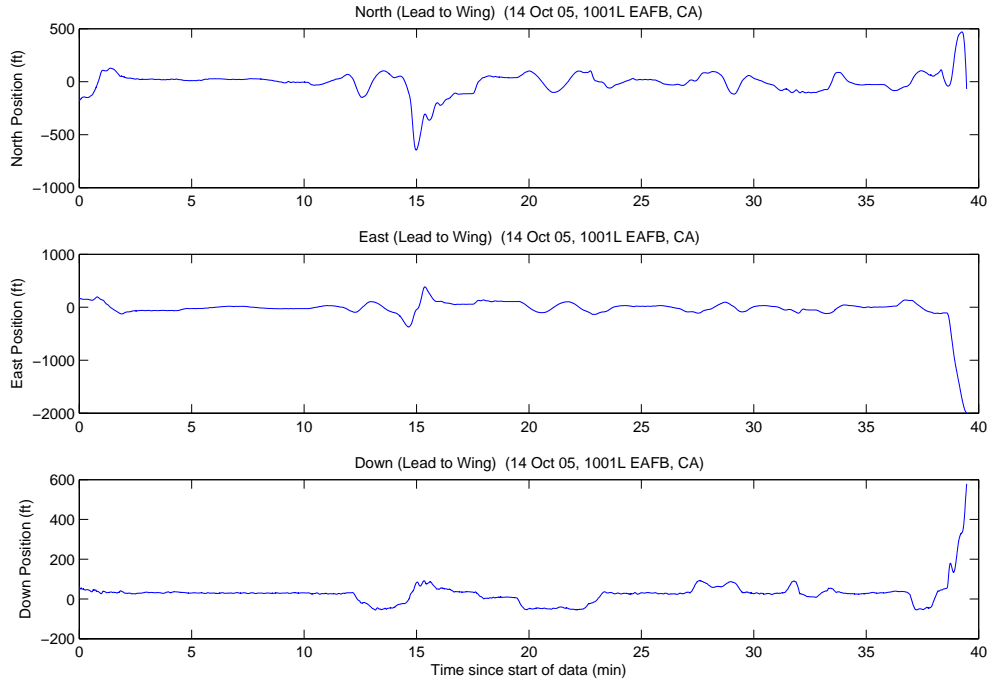


Figure A.382: Case F8.2: North, East, Down Relative Positions

A total of 30 unique ambiguity sets were returned from the LAMBDA function during the data run. The aircraft were 77.9 meters apart when the system started, with eight visible satellites. The true ambiguity set was the first set returned from the LAMBDA function. The minimum indicator correctly chose the true ambiguity set at four minutes, and did not change throughout the rest of the data run. The true set was returned from LAMBDA as number one, 2247 of 2247 epochs (100.0 percent), with a maximum ratio of 311.9.

Figures A.383, A.384, and A.385 show the floating filter position errors. The floating filter performed well, and was very similar to the previous data runs. The floating filter's 1σ estimate of the error is outside the plot scales in all three figures.

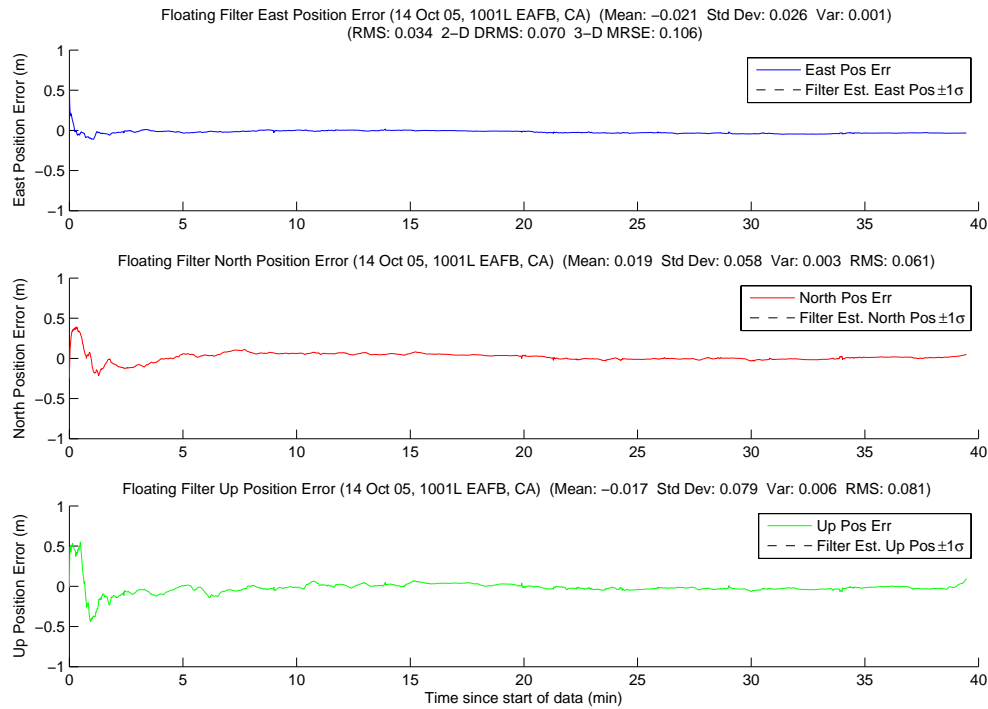


Figure A.383: Case F8.2: Floating Filter Position Errors

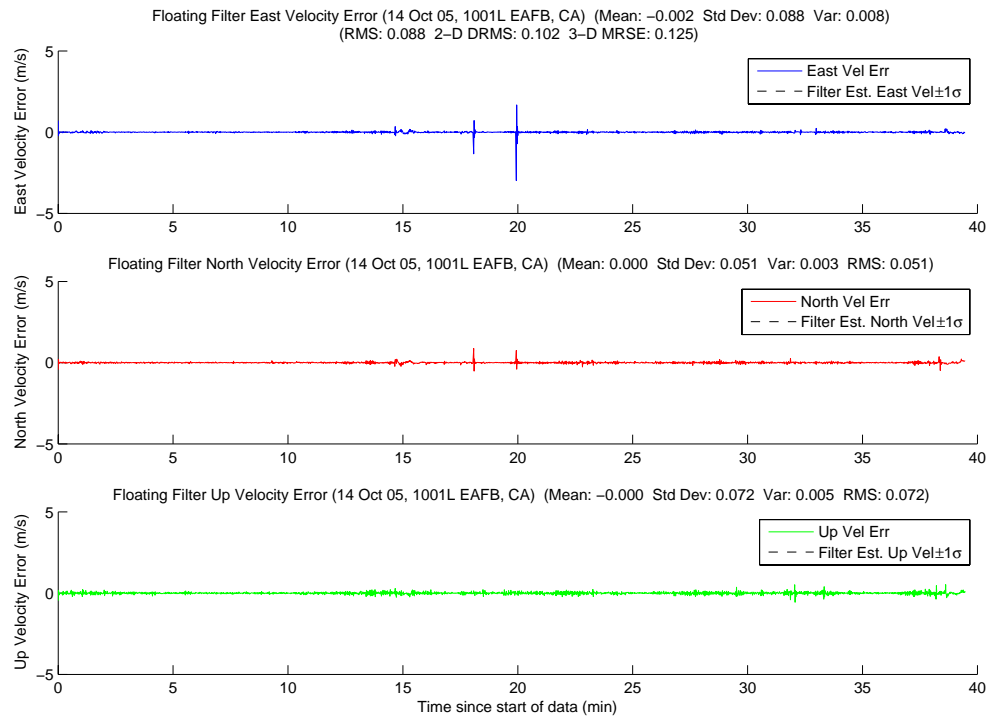


Figure A.384: Case F8.2: Floating Filter Velocity Errors

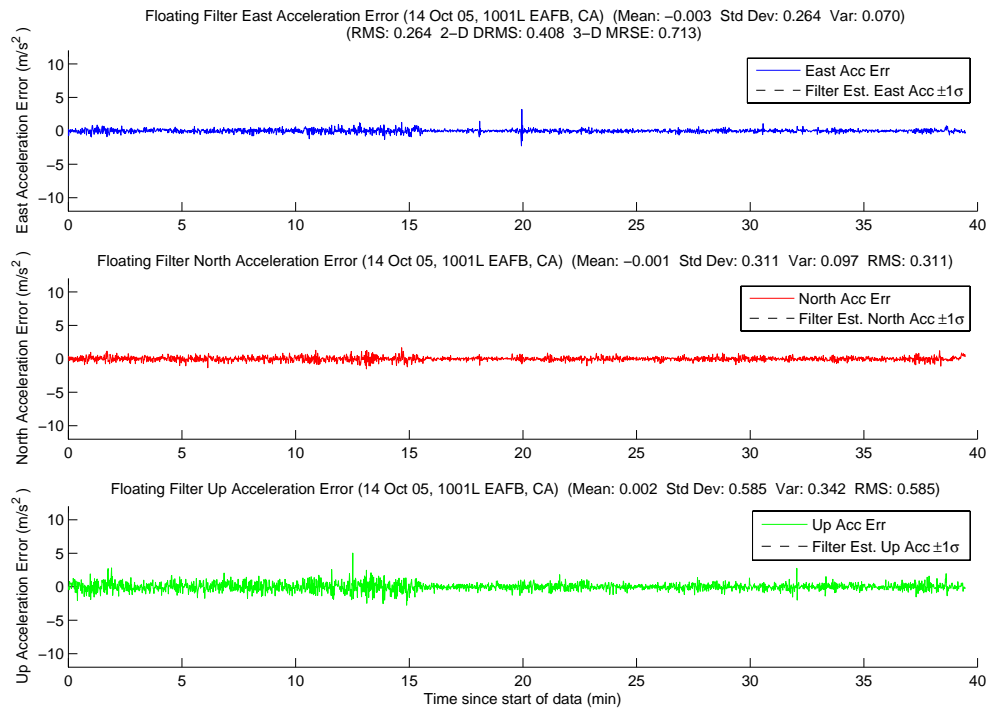


Figure A.385: Case F8.2: Floating Filter Acceleration Errors

Figure A.386 displays PRN 8 measurement information and the floating filter ambiguity estimate. PRN 8 masked nine times and had one cycle slip during the run. The cycle slip occurred at 14 minutes. The rest of the satellite's plots follow in Figures A.387, A.388, A.389, A.390, A.391, A.392, and A.393.

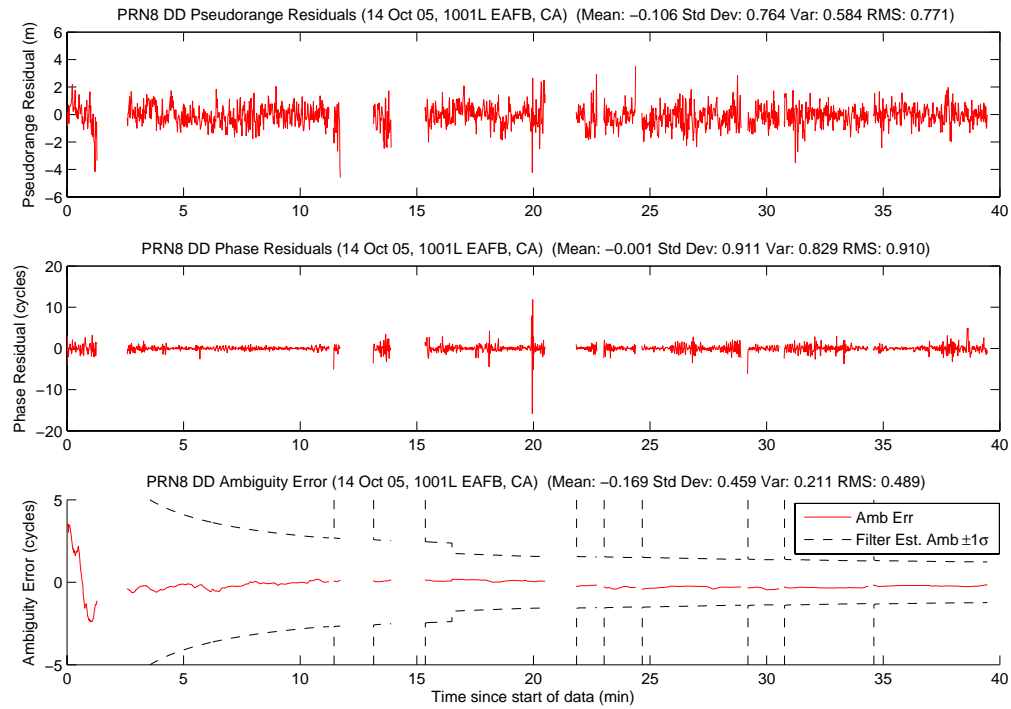


Figure A.386: Case F8.2: Satellite 8 Measurements

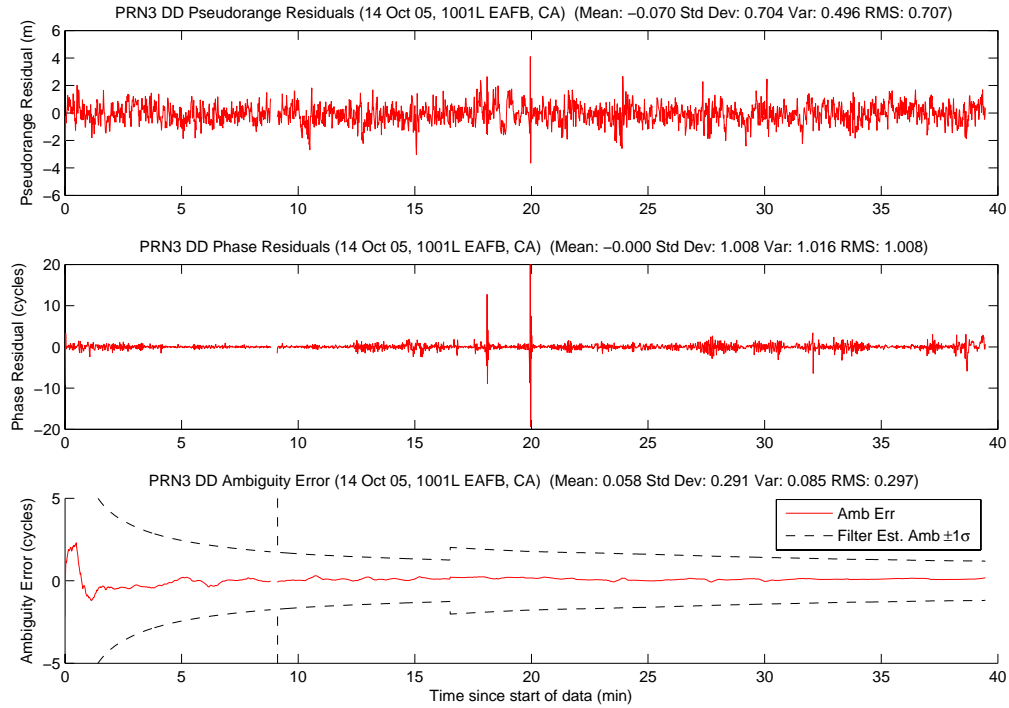


Figure A.387: Case F8.2: Satellite 3 Measurements

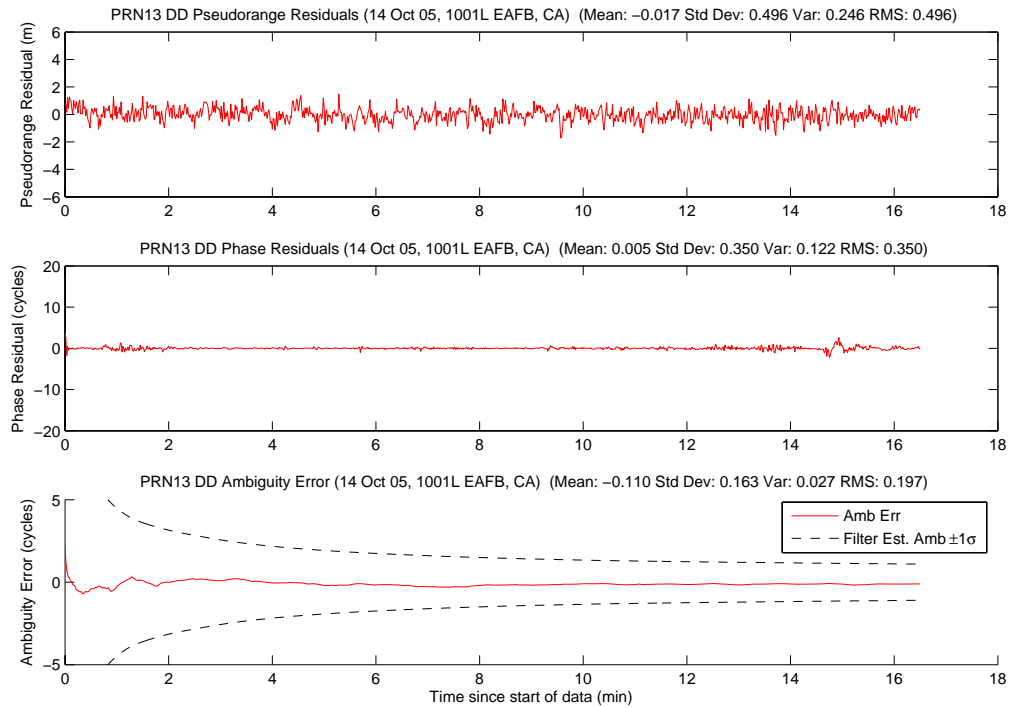


Figure A.388: Case F8.2: Satellite 13 Measurements

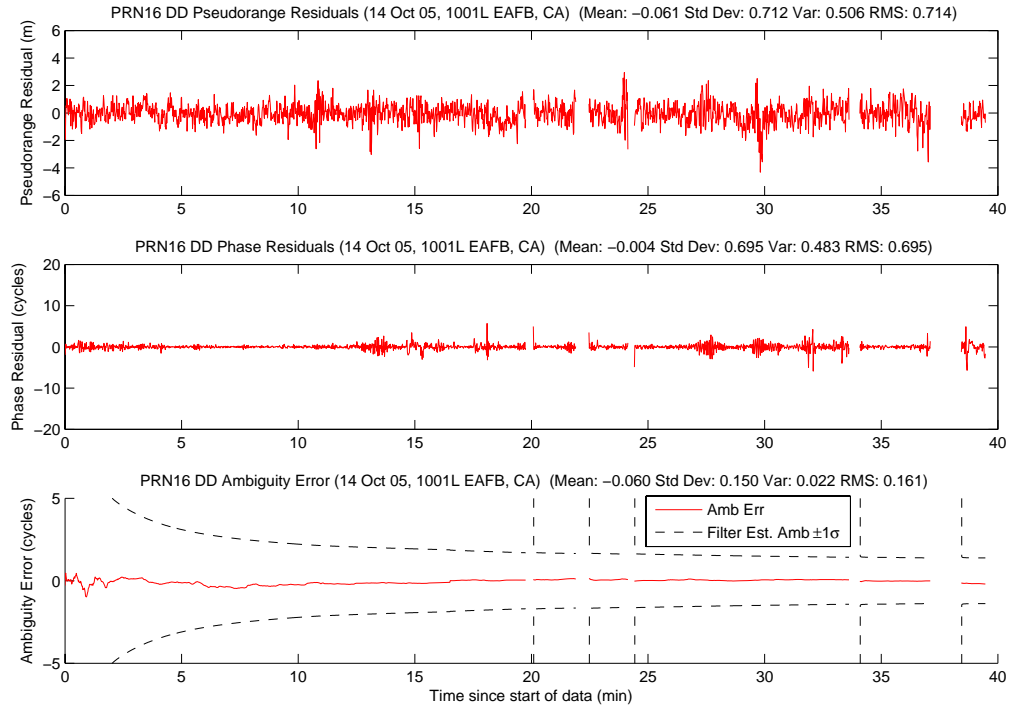


Figure A.389: Case F8.2: Satellite 16 Measurements

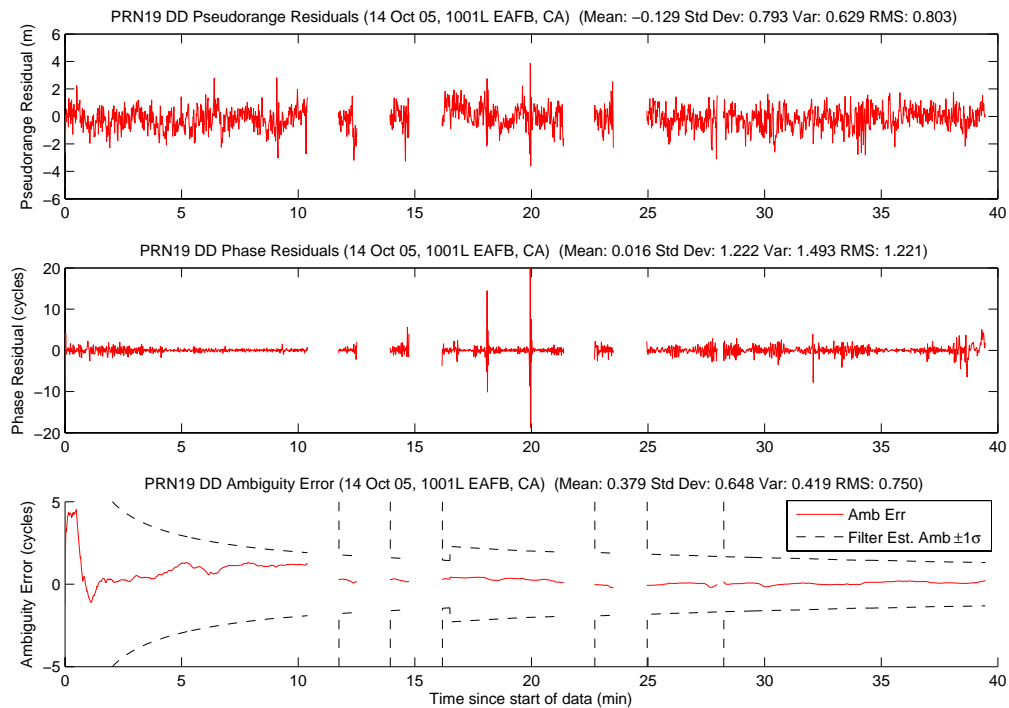


Figure A.390: Case F8.2: Satellite 19 Measurements

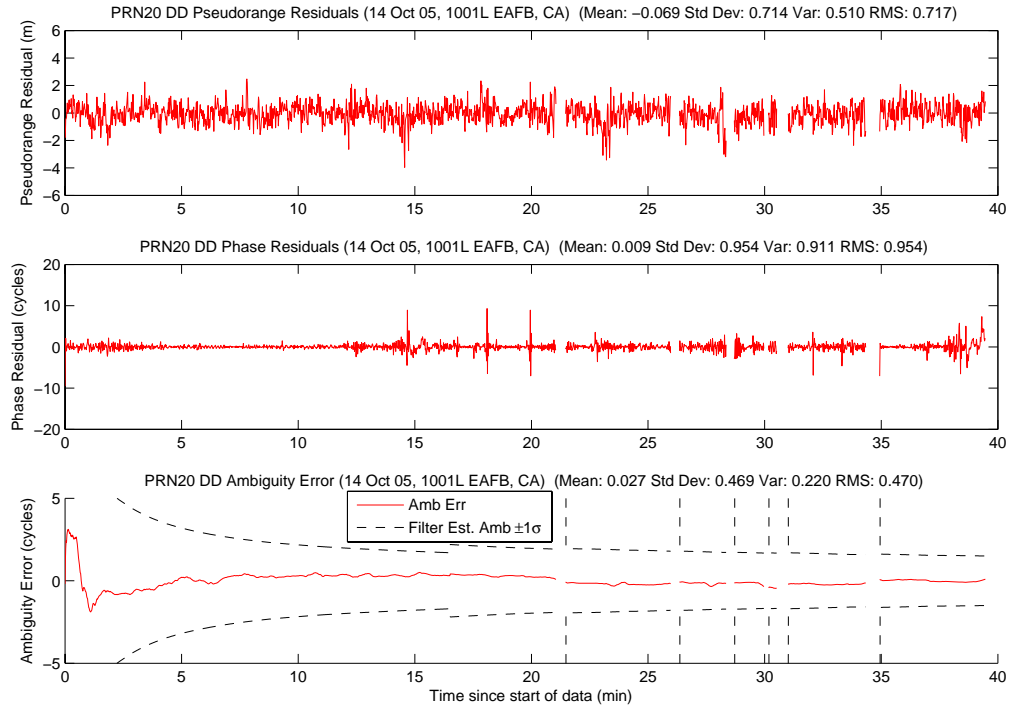


Figure A.391: Case F8.2: Satellite 20 Measurements

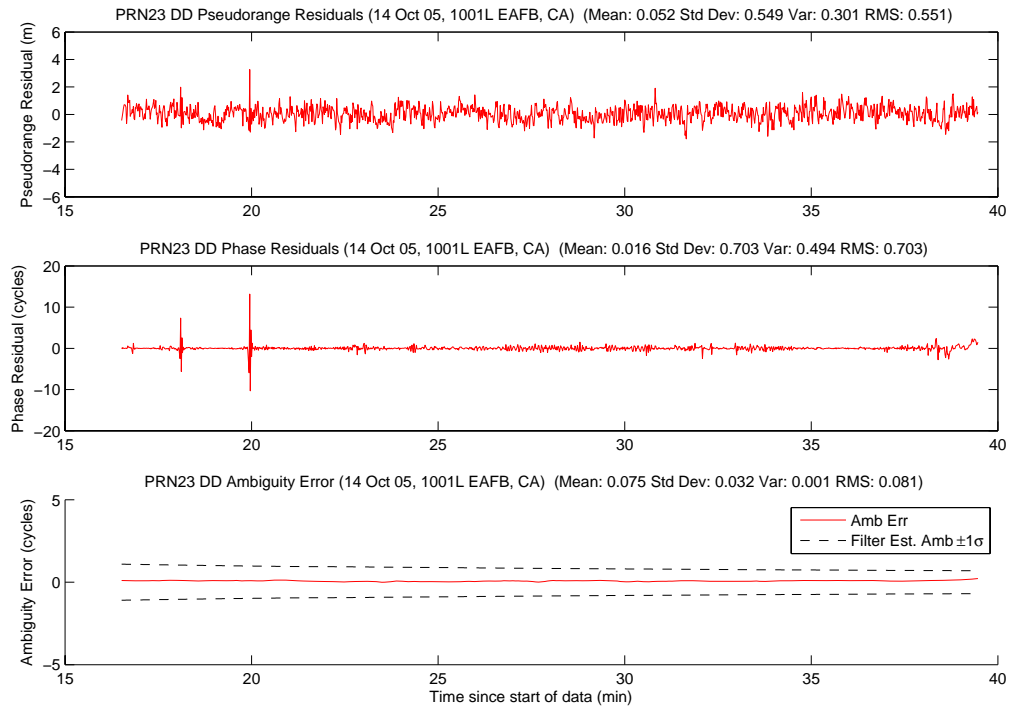


Figure A.392: Case F8.2: Satellite 23 Measurements

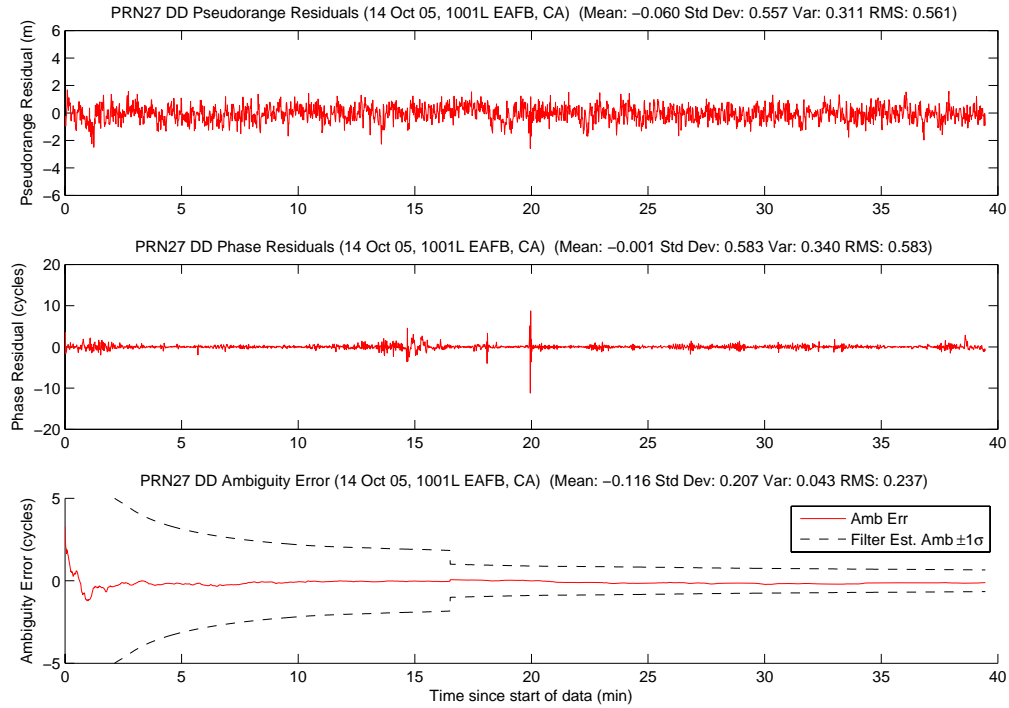


Figure A.393: Case F8.2: Satellite 27 Measurements

Figure A.394 shows the number of visible satellites. There was substantial satellite activity during the data run. There were 28 dropouts, with no satellites rising or setting. Of the 28 dropouts, three were due to cycle slips. Each cycle slip was correctly identified by the triple-difference method discussed in Section 3.9.1. The cycle slips occurred on PRNs 3, 8, and 20 at 9, 14, and 26 minutes, respectively. PRN 8 had the most dropouts with 10. The rest were split fairly evenly between PRNs 16, 19, and 20. This was the most dropouts and cycle slips of any data run. There was a third aircraft in the formation taking video, which was the most likely cause for the increase in dropouts and cycle slips. The filter, MMAE, and minimum indicator still performed very well, a testament to the algorithm's robustness.

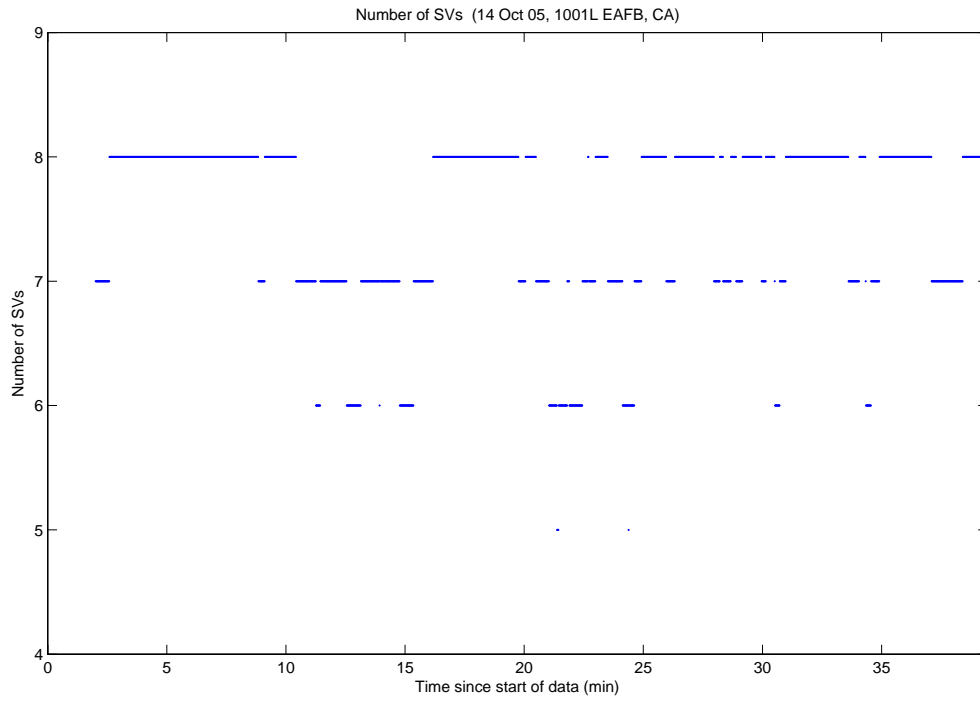


Figure A.394: Case F8.2: Visible Satellites

Figure A.395 shows all of the candidate ambiguity set's SOS residuals. The residual plot is very similar to the last data run. Nothing in particular stands out. Figure A.396 displays the MMAE conditional probability. The probability plot was a little surprising, especially considering the number of satellite dropouts. However, the MMAE was able to correctly identify the true ambiguity set for the entire run.

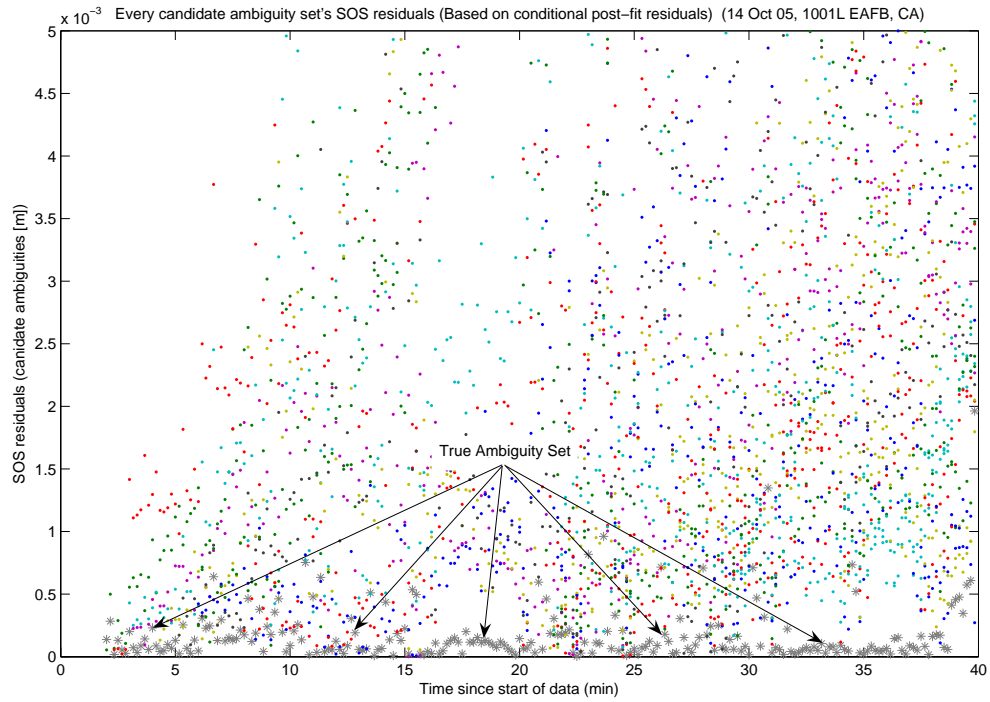


Figure A.395: Case F8.2: SOS Residuals

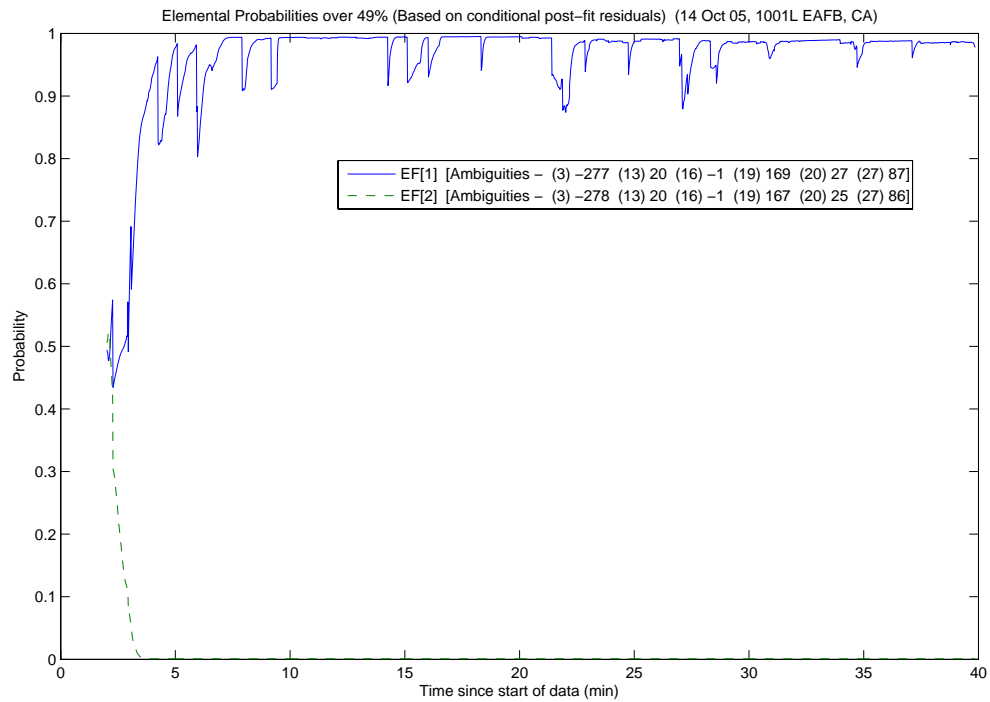


Figure A.396: Case F8.2: Conditional Probabilities for Selected MMAE Elemental Filters

The following figures show the difference between the “pre” and “post-fit” residuals in the conditional probability calculation. Figure A.397 is EF[1], and Figure A.398 is EF[2].

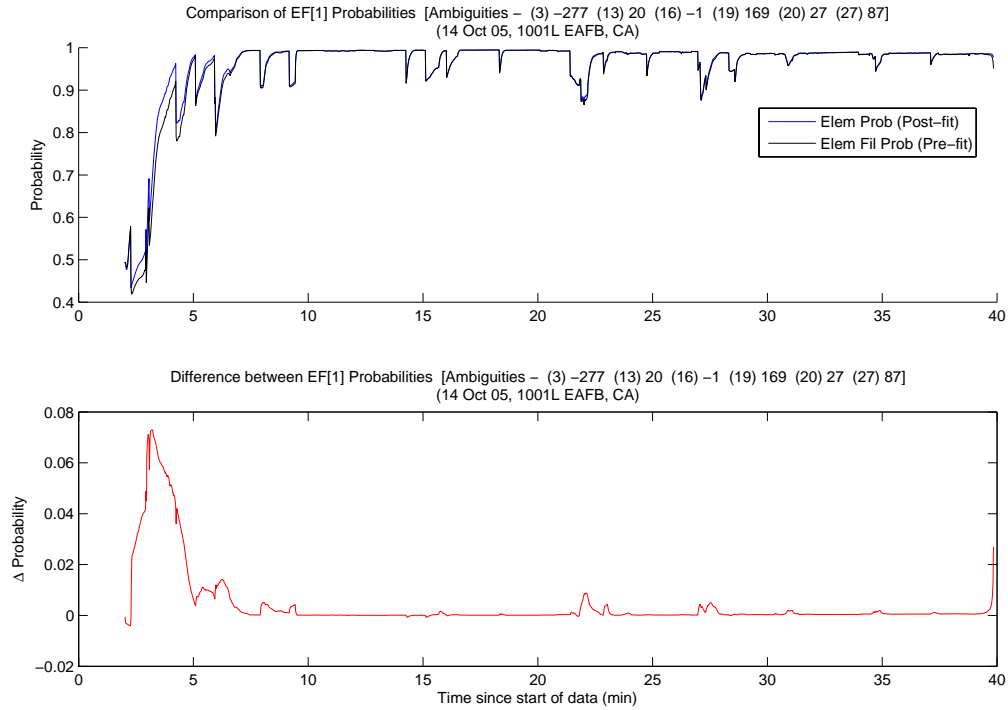


Figure A.397: Case F8.2: EF[1] Probability Comparison

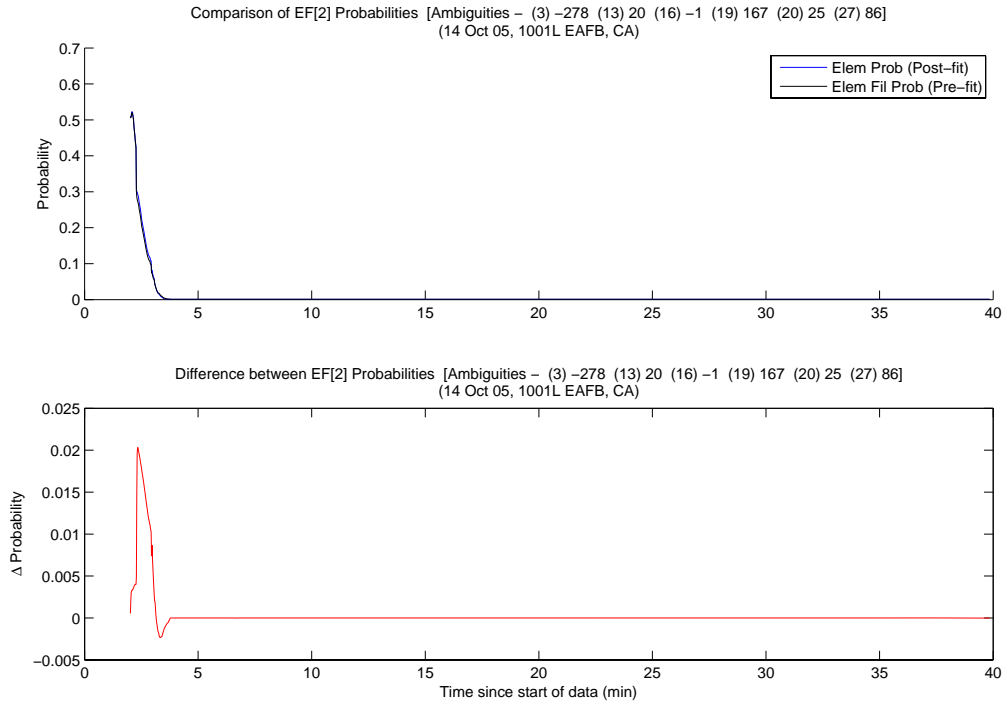


Figure A.398: Case F8.2: EF[2] Probability Comparison

The MMAE position error is shown in Figure A.399.

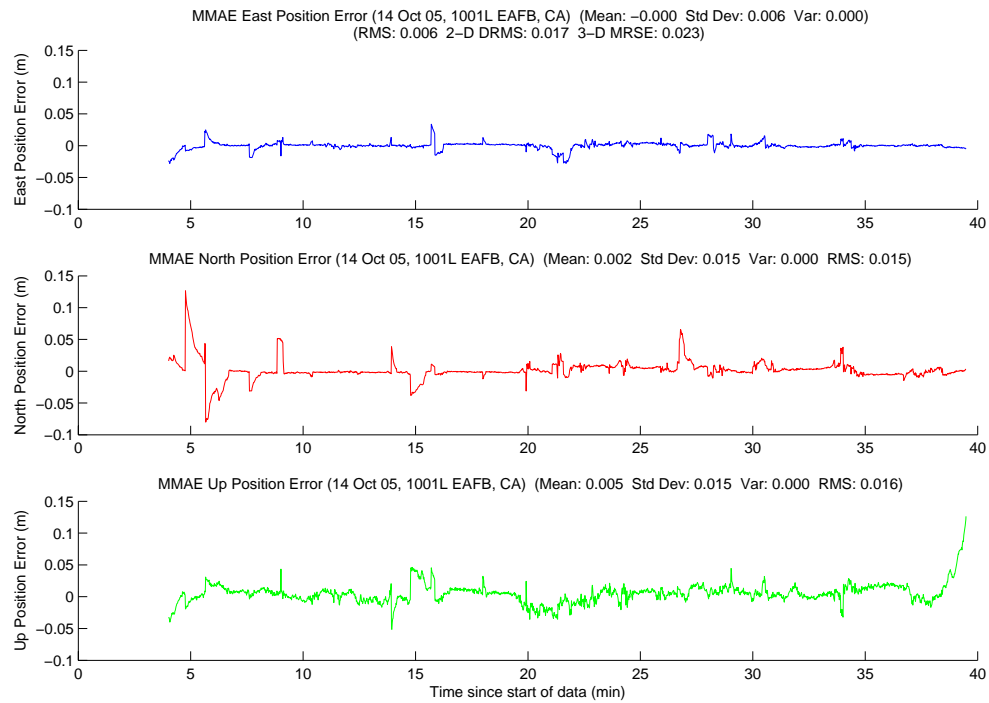


Figure A.399: Case F8.2: MMAE Position Error

Finally, Figure A.400 displays the position error plot for the minimum indicator variable. The position errors on the minimum indicator had numerous small jittery spikes, especially from 10 to 15 minutes, and from 20 to 35 minutes. These times corresponded to when there were numerous dropouts and cycle slips. Phase measurements from the Javad[®] receivers were noisy (possibly due to aircraft masking of satellites and potential multipath). The minimum indicator MRSE was still well within the centimeter-level requirement at 0.016 meters. Lastly, Table A.21 gives a summary of each method's position error.

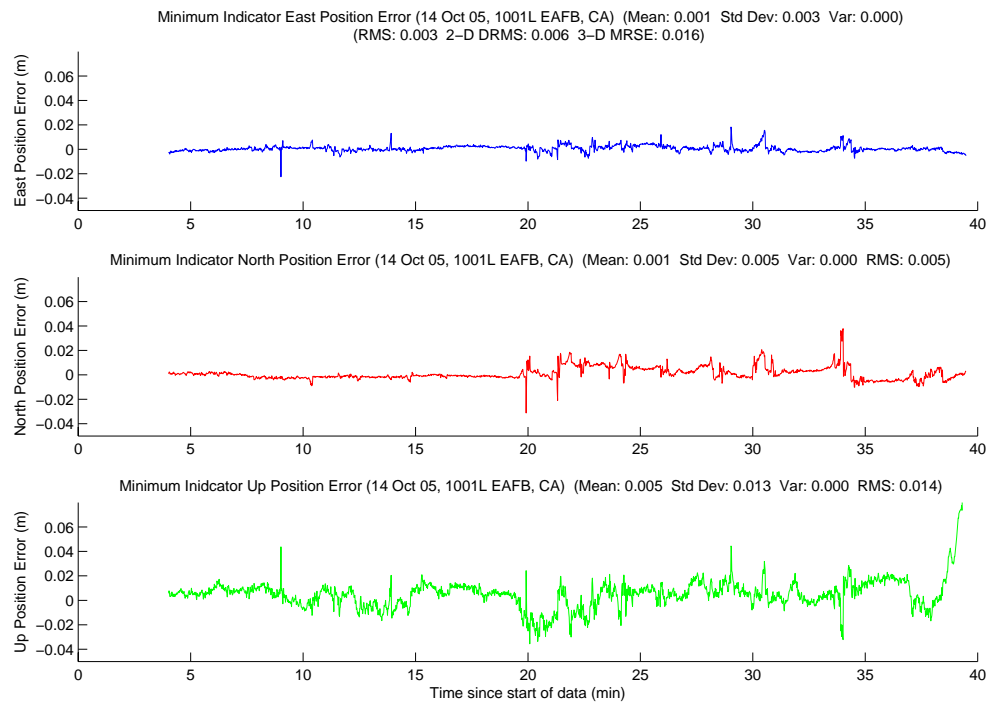


Figure A.400: Case F8.2: Minimum Indicator Position Error

Table A.21: Case F8.2: Position Error Summary (m)

	East		North		Up		DRMS	MRSE
	Error	Std	Error	Std	Error	Std		
Float Filter	-0.021	0.026	0.019	0.058	-0.017	0.079	0.070	0.106
MMAE	-0.000	0.006	0.002	0.015	0.005	0.015	0.017	0.023
Min. Ind.	0.001	0.003	0.001	0.005	0.005	0.013	0.006	0.016

Bibliography

1. Bastos, L. and H. Landau. "Fixing Cycle Slips in Dual-Frequency Kinematic GPS-Applications Using Kalman Filtering". *Manuscripta Geodaetica*, 13(4):249–256, 1988.
2. Bisnath, Sunil B. "Efficient, Automated Cycle-Slip Correction of Dual-Frequency Kinematic GPS Data". *The Satellite Division of the Institute of Navigation's 13th International Technical Meeting*, 145–154. Salt Lake City, UT, 19-22 September 2000.
3. Blewitt, G. "An Automatic Editing Algorithm for GPS Data". *Geophysical Research Letters*, 17(3):199–202, 1990.
4. Bouska, Terry J. *Development and Simulation of a Pseudolite-Based Flight Reference System*. MS thesis, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB, OH, March 2003. AFIT/GE/ENG/03-03.
5. Bowers, Roshawn E. *Estimation Algorithm for Autonomous Aerial Refueling Using a Vision Based Relative Navigation System*. MS thesis, Texas A&M University, College Station, TX, August 2005.
6. Chen, D. "Fast Ambiguity Search Filter (FASF): A Novel Concept for GPS Ambiguity Resolution". *Proceedings of ION GPS-93, 6th International Technical Meeting of the Satellite Division of the Institute of Navigation*, 781–787. Salt Lake City, UT, 22-24 September 1993.
7. Chen, D. and Gérard Lachapelle. "A Comparison of the FASF and Least-Squares Search Algorithms for On-the-Fly Ambiguity Resolution". *NAVIGATION: Journal of the Institute of Navigation*, 42(2):371–390, Summer 1995.
8. Counselman, C.C. and S.A. Gourevitch. "Miniature Interferometer Terminals for Earth Surveying: Ambiguity and Multipath for the Global Positioning System". *IEEE Transactions on Geoscience and Remote Sensing*, GE-19(4):244–252, October 1981.
9. Crow, S.C. and F.L. Manning. "Differential GPS Control of Starcar 2". *NAVIGATION: Journal of the Institute of Navigation*, 39(4):383–405, Winter 1992.
10. Davis, Brooke. "First-ever Coordinated UAV Flight Makes History". http://www.edwards.af.mil/archive/2004/2004-archive-2_x45a.html, August 2004.
11. Diamond Systems, Inc. "PB-EAP-500-K". <http://www.diamondsystems.com/products/pandora>, 2001-2006. Mountain View, CA.

12. Edwards, F.G. and D.M. Hegarty. "Flight-test Evaluation of Civil Helicopter Terminal Approach Operations Using Differential GPS". *AIAA Guidance, Navigation and Control Conference*, 657–666. Boston, MA, 14-16 August 1989.
13. Fulghum, David. "Hands-Off Aerial Refueling". *Aviation Week & Space Technology*, 62, 12 December 2005.
14. Gao, Y. and Z. Li. "Cycle Slip Detection and Ambiguity Resolution Algorithms for Dual-Frequency GPS Data Processing". *Marine Geodesy*, 22(4):169–181, 1999.
15. Gehue, H., G. Lachapelle, M.E. Cannon, T.W. Goddard, and D.C. Penney. "GPS System Integration and Field Approaches in Precision Farming". *Proceedings of National Technical Meeting of the Satellite Division of the Institute of Navigation*, 549–556. San Diego, CA, 24-26 January 1994.
16. George, Benjamin, Bruce Wilder, York Pasanen, Adam Faulkner, and Scott Sullivan. *Limited Evaluation of a Relative GPS Datalink Between Two C-12C Aircraft*. Technical Report AFFTC-TIM-05-04, USAF, 220 South Wolfe Avenue, Edwards AFB, CA 93524-6485, June 2005.
17. Goad, C. "Precise Positioning with the Global Positioning System". *Proceedings of the Third International Symposium on Inertial Technology for Surveying and Geodesy*, 745–756. Banff, Canada, 16-20 September 1985.
18. Hanson, Curtis E., Jack Ryan, Michael J. Allen, and Steven R. Jacobson. *An Overview of Flight Test Results for a Formation Flight Autopilot*. Technical Report NASA/TM-2002-210729, NASA Dryden Flight Research Center, Edwards AFB, CA, August 2002.
19. Hatch, R. "Instantaneous Ambiguity Resolution". *Proceedings of KIS-90, The 1990 International Symposium on Kinematic Systems in Geodesy, Geomatics, and Navigation*, 299–308. Banff, Canada, 10-13 September 1990.
20. Henderson, Paul E. *Development and Testing of a Multiple Filter Approach for Precise DGPS Positioning and Carrier-Phase Ambiguity Resolution*. MS thesis, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB, OH, March 2001. AFIT/GE/ENG/01M-15.
21. Hundley, Warren, Stephen Rowson, and Glenn Courtney. "Flight Evaluation of a Basic C/A Code Differential GPS for Category I Precision Approach". Wilcox Electric, Inc.
22. Joffrion, Jacques M. *Head Tracking for 3D Audio Using a GPS-Aided MEMS IMU*. MS thesis, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB, OH, March 2005. AFIT/GE/ENG/05-09.
23. Jonge de, Paul and Christian Tiberius. *The LAMBDA Method for Integer Ambiguity Estimation: Implementation Aspects*. LGR-Series No. 12, Delft University of Technology, Faculty of Geodetic Engineering, Thijsseweg 11, 2629 JA Delft, The Netherlands, August 1996.

24. Jonge de, P.J., C.C.J.M. Tiberius, and P.J.G. Teunissen. "Computational Aspects of the LAMBDA Method for GPS Ambiguity Resolution". *Proceedings of ION GPS-96, 9th International Technical Meeting of the Satellite Division of the Institute of Navigation*, 935–944. Kansas City, Missouri, 17-20 September 1996.
25. Joosten, Peter. *The LAMBDA-Method: Matlab™ Implementation*. In-house, Delft University of Technology, Faculty of Geodetic Engineering, Thijsseweg 11, 2629 JA Delft, The Netherlands, 12 March 2001.
26. Joosten, Peter. "LAMBDA Method". <http://www.lr.tudelft.nl/live/-pagina.jsp?id=e9184f4b-9751-417c-bb3d-239af1deaf4a&lang=en>, 2002-2004. Delft, Netherlands.
27. Joosten, Peter and Christian Tiberius. "Fixing the Ambiguities: Are You Sure They're Right?" *GPS World*, 11(5):46–51, May 2000.
28. KDE. "KDevelop 3.2.0". <http://www.kdevelop.org>, 2004-2005.
29. Klobuchar, J.A. *Global Positioning System: Theory and Applications, Volume I*, chapter 12: Ionospheric Effects in GPS, 485–515. American Institute of Aeronautics and Astronautics, Inc., Washington, 1996.
30. Lachapelle, Gérard, M.E. Cannon, H. Gehue, T.W. Goddard, and D.C. Penny. "GPS Systems Integration and Field Approaches in Precision Farming". *NAVIGATION: Journal of the Institute of Navigation*, 41(3):323–335, Fall 1994.
31. Mader, G.L. "Ambiguity Function Techniques for GPS Phase Initialization and Kinematic Solution". *Proceedings of the 2nd International Symposium on Precise Positioning with GPS*, 1234–1247. Canadian Institute of Surveying and Mapping, Ottawa, Canada, September 1990.
32. Maybeck, Peter S. *Stochastic Models, Estimation, and Control, I*. New York: Academic Press Inc., 1979. Republished, Arlington, VA: Navtech, 1994.
33. Maybeck, Peter S. *Stochastic Models, Estimation, and Control, II*. New York: Academic Press Inc., 1982. Republished, Arlington, VA: Navtech, 1994.
34. 412 TW/TSRE. "GPS Aided Inertial Navigation Reference (GAINR) System". Edwards AFB, California, 1 August 1995.
35. Javad Navigation Systems, Inc. "JNS100". <http://www.javad.com>, 2004. Saratoga, CA.
36. NovAtel Inc. "GrafMov 7.50.2426". <http://www.novatel.com/products/-waypoint-grafmov.htm>, 2005. Calgary, Alberta, Canada.
37. The MathWorks, Inc. "MATLAB® Release 2006a". <http://www.mathworks.com/>, 1994-2006. Natick, MA.
38. Xsens Technologies B.V. "MT9-B". <http://www.xsens.com/index.php?-mainmenu=products&submenu=general&subsubmenu=MT9-B>, 2006. Enschede, Netherlands.

39. MicroBee Systems, Inc. "PCFW-104". <http://www.microbee-systems.com/-PCFW.htm>, 2002-2006. Colbert, WA.
40. Microbotics, Inc. "MIDG II INS/GPS". http://microboticsinc.com/-ins_gps.php, 2006. Hampton, VA.
41. Misra, Pratap and Per Enge. *Global Positioning System: Signals, Measurements, and Performance*. Ganga-Jamuna Press, Lincoln, MA, 2004.
42. Novell, Inc. "SUSE 9.2". <http://www.novell.com/linux>, 2004-2005. Waltham, MA.
43. O'Connor, Michael L., Gabriel H. Elkaim, and Bradford W. Parkinson. "Carrier-Phase DGPS for Closed-Loop Control of Farm and Construction Vehicles". *NAVIGATION: Journal of the Institute of Navigation*, 43(2):167–178, Summer 1996.
44. Ormsby, Charles D. *Generalized Residual Multiple Model Adaptive Estimation of Parameters and States*. Ph.D. dissertation, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB, OH, October 2003. AFIT/DS/ENG/03-08.
45. Osterroos, Ryan K. *Full Capability Formation Flight Control*. MS thesis, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB, OH, February 2005. AFIT/GAE/ENY/05-M16.
46. Pointon, J. and K. Babu. "LANDNAV: A Highly Accurate Land Navigation System for Agricultural Applications". *Proceedings of ION GPS-94, 7th International Technical Meeting of the Satellite Division of the Institute of Navigation*, 1077–1080. Salt Lake City, UT, 20-23 September 1994.
47. Raquet, John F. "Class Notes: EENG633 GPS Receiver Design". Summer Quarter, 2004, Air Force Institute of Technology, WPAFB, OH.
48. Raquet, John F. *Development of a Method for Kinematic GPS Carrier-Phase Ambiguity Resolution Using Multiple Reference Receivers*. Ph.D. dissertation, UCGE Reports Number 20116, Department of Geomatics Engineering, University of Calgary, Calgary, Alberta, Canada, May 1998.
49. Remondi, B.W. "Pseudo-Kinematic GPS Results using the Ambiguity Function Method". *NAVIGATION: Journal of the Institute of Navigation*, 38(1):17–36, Spring 1991.
50. Ross, Steven, Elwood Waddell Jr., Aaron Mainstone, and Juanluis Velez. *Demonstration of a Control Algorithm for Autonomous Aerial Refueling (AAR)*. Technical Report AFFTC-TIM-05-10, USAF, 220 South Wolfe Avenue, Edwards AFB, CA 93524-6485, December 2005.
51. Ross, Steven M. *Formation Flight Control for Aerial Refueling*. MS thesis, Graduate School of Engineering, Air Force Institute of Technology (AETC), Wright-Patterson AFB, OH, March 2006. AFIT/GAE/ENY/06-M35.

52. Schlechte, Gene L. *Design Process for the United States Coast Guard's Differential GPS Navigation Service*. Technical report, U.S. Coast Guard Omega Navigation System Center, 7323 Telegraph Road, Alexandria, VA 22310-3998, 1994.
53. Teunissen, P.J.G. "A New Method for Fast Carrier Phase Ambiguity Estimation". *Proceedings of IEEE PLANS-94, IEEE Position, Location, and Navigation Symposium*, 562–573. Las Vegas, Nevada, 11-15 April 1994.
54. Teunissen, P.J.G. "The Least-Squares Ambiguity Decorrelation Adjustment: A Method for Fast GPS Integer Ambiguity Estimation". *Journal of Geodesy*, 70(1-2):65–82, November 1995.
55. Van Dierendonck, A.J. "Understanding GPS Receiver Terminology: A Tutorial on What Those Words Mean". *Proceedings of KIS-94, The 1994 International Symposium on Kinematic Systems in Geodesy, Geomatics, and Navigation*, 15–24. Banff, Alberta, Canada, 30 August - 2 September 1994.
56. Wei, M. and K.P. Schwarz. "Fast Ambiguity Resolution Using an Integer Non-linear Programming Method". *Proceedings of ION GPS-95, 8th International Technical Meeting of the Satellite Division of the Institute of Navigation*, 1101–1110. Palm Springs, CA, 12-15 September 1995.

Vita

Major Christopher J. Spinelli was born in Abilene, Texas. He grew up in Richardson, Texas where he graduated from L.V. Berkner High School in 1992. He then attended the United States Air Force Academy (USAFA) in Colorado Springs, Colorado where he earned his Bachelor of Science Degree in Computer Science in 1996.

Having been appointed a Second Lieutenant in the United States Air Force in 1996, his first assignment was Undergraduate Pilot Training (UPT) at Vance AFB, OK. He was selected for F-16 replacement training in the 63rd Fighter Squadron at Luke AFB, AZ following pilot training. In October of 1998, Major Spinelli was transferred to Moody AFB, GA to fly the F-16CG operationally in the 69th Fighter Squadron. While at Moody AFB, he went on numerous deployments including combat operations over Iraq for OPERATION SOUTHERN WATCH (OSW). Two years later in 2000, Major Spinelli transferred again, now to Hill AFB, UT as part of the 421st Fighter Squadron. During his assignment at Hill AFB, Major Spinelli upgraded to Instructor Pilot (IP) in the F-16CG and was a Flight Commander for the last year of his assignment. Also, in 2003, he deployed to Qatar for OPERATION IRAQI FREEDOM (OIF).

In 2004, he was selected to attend the joint Air Force Institute of Technology (AFIT) and USAF Test Pilot School (TPS) Program. Maj Spinelli graduated from TPS as an Experimental Test Pilot in June 2006, and will receive his Masters Degree in Computer Science in September 2006. Following AFIT/TPS, his assignment is to the 416th Flight Test Squadron at Edwards AFB, CA where he will be flight-testing the F-16.

Permanent address: ON FILE

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188							
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.												
1. REPORT DATE (DD-MM-YYYY) 14-09-2006		2. REPORT TYPE Master's Thesis			3. DATES COVERED (From — To) Mar 2004 — Sept 2006							
4. TITLE AND SUBTITLE <div style="text-align: center;">Development and Testing of a High-Speed Real-Time Kinematic Precise DGPS Positioning System Between Two Aircraft</div>					5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER 5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER 							
6. AUTHOR(S) Spinelli, Christopher, J., Major, USAF					8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GCS/ENG/06-12							
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way, Building 640 WPAFB OH 45433-7765					10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S) 							
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Mikel M. Miller, Ph.D. Sensors Directorate, Air Force Research Laboratory (AFRL/SNRN) Bldg 620, Room 3AJ39 2241 Avionics Circle WPAFB, OH 45433-7333 (937) 255-6127 x4274 mikel.miller@wpafb.af.mil					12. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE, DISTRIBUTION UNLIMITED							
13. SUPPLEMENTARY NOTES												
14. ABSTRACT This research involves the design, implementation, and testing of a high-speed, real-time kinematic, precise differential GPS positioning system for use in airborne applications such as automated aerial-refueling and close formation flying. Although many of the current ambiguity resolution techniques use the residuals from the least squares position estimation to determine the true ambiguity set, this thesis presents a novel approach to the ambiguity resolution problem, called the minimum indicator. Instead of assuming the ambiguity set with the lowest residuals is the true set, other special characteristics of the residuals are examined. This increases the confidence that the algorithm has selected the true ambiguity set. The end result was the first-ever successful in-flight demonstration of close formation flight, culminating in over 11 hours of close formation flying with a mean radial spherical error of 3.3 centimeters (0.108 feet). Other areas addressed include: the difference between “pre-fit” and “post-fit” residuals in the conditional probability calculation, the impact of a simplified dynamics model on system performance, the effect of widelane observables on the time-to-fix the correct double-difference ambiguities, and dynamically adjusting the time constant and standard deviation of relative acceleration states in the FOGMA model.												
15. SUBJECT TERMS Differential GPS (DGPS), Relative, Precise, Widelane, Kalman Filter, Minimum Indicator, Multiple Model Adaptive Estimation (MMAE), Real-time Kinematic (RTK), LAMBDA, Microelectromechanical Systems (MEMS), Inertial Measurement Unit (IMU), C-12 Aircraft, Learjet, Flight Testing, Two Aircraft, Aerial-refueling, Autonomous												
16. SECURITY CLASSIFICATION OF: <table border="1" style="width: 100%; border-collapse: collapse; font-size: x-small;"> <tr> <td style="width: 33%; text-align: center;">a. REPORT</td> <td style="width: 33%; text-align: center;">b. ABSTRACT</td> <td style="width: 33%; text-align: center;">c. THIS PAGE</td> </tr> <tr> <td style="text-align: center;">U</td> <td style="text-align: center;">U</td> <td style="text-align: center;">U</td> </tr> </table>			a. REPORT	b. ABSTRACT	c. THIS PAGE	U	U	U	17. LIMITATION OF ABSTRACT <div style="text-align: center;">UU</div>		18. NUMBER OF PAGES <div style="text-align: center;">483</div>	
a. REPORT	b. ABSTRACT	c. THIS PAGE										
U	U	U										
19a. NAME OF RESPONSIBLE PERSON Dr. John F. Raquet AFIT/ENG			19b. TELEPHONE NUMBER (include area code) (937) 255-3636, ext 4580									